arXiv:1512.08815v1 [stat.ME] 29 Dec 2015

A Pivot-Based Improvement to Sandwich-Based Confidence Intervals

James W. Harmon[1] and Peter D. Hoff[1,2]

Departments of Statistics[1] and Biostatistics[2]

University of Washington

December 31, 2015

Abstract

The current standard for confidence interval construction in the context of a possibly misspecified model is to use an interval based on the sandwich estimate of variance. These intervals provide asymptotically correct coverage, but small-sample coverage is known to be poor. By eliminating a plug-in assumption, we derive a pivot-based method for confidence interval construction under possibly misspecified models. When compared against confidence intervals generated by the sandwich estimate of variance, this method provides more accurate coverage of the pseudo-true parameter at small sample sizes. This is shown in the results of several simulation studies. Asymptotic results show that our pivot-based intervals have large sample efficiency equal to that of intervals based on the sandwich estimate of variance.

1 Introduction

The method of maximum likelihood is a powerful and efficient technique for statistical inference, but much of the efficiency is lost if the data-generating model is incorrectly specified. If a statistician selects the wrong family of distributions to model the data, then the standard variance estimates using the Fisher information matrix will likely be biased, even asymptotically. Consequently, any inference performed with the incorrect model will be adversely affected. An example of this is when

true coverage rates of confidence intervals do not match nominal coverage rates even as the sample size increases.

This begs the question of why any analyst would want to use maximum likelihood or any parametric methods at all. One answer is that the parameter in a misspecified model might still be the appropriate parameter to estimate. An example of this might be the measure of linear association between a covariate and the response variable in a regression model. Even if the full model has been misspecified, this association might still be of interest to a researcher. Given certain regularity conditions, the maximum likelihood estimator converges to a unique pseudo-true parameter. This pseudo-true parameter has a nice interpretation: it minimizes the Kullback-Leibler distance between the true data-generating model and assumed family of models [Akaike, 1973]. In other words, the pseudo-true parameter yields the distribution from the misspecified model that is "closest" to the true data-generating distribution. The pseudo-true parameter might well be of interest, even if the model is misspecified. For example, an analyst might be interested in knowing the linear association between a covariate and the response in a linear regression.

The second answer is that parametric methods increase inferential efficiency, even if the model is incorrect. As one moves through the spectrum of parametric, semi-parametric, and non-parametric methods, assumptions about the data-generating model decrease. The trade-off for fewer assumptions is that the efficiency of one's inference decreases as well. In terms of confidence intervals, decreased efficiency means larger confidence intervals for a given confidence level and sample size. Using fully empirical methods clearly requires the fewest assumptions of the true data-generating model, but the cost in efficiency might be too high. Once the decision to use parametric or semi-parametric methods has been made, the question then boils down to one of how to perform inference about the pseudo-true parameter.

The standard method for creating confidence intervals using misspecified models was introduced by Huber [1967] and expanded by White [1982]. It is based on what is commonly called the "sandwich" estimate of variance, so named because the matrix version of this estimator has the form $A^{-1}BA^{-1}$ where the outer matrices are the "bread" and the inner matrix is the "peanut butter". The main benefit of the sandwich estimate of variance is that it provides asymptotically

correct variance estimates for maximum likelihood parameter estimates. This yields confidence intervals with asymptotically correct coverage rates for pseudo-true parameters. Additionally, it is applicable across a wide range of statistical techniques and is not limited to maximum likelihood methods [Chen and Fan, 2005], [Huber and Ronchetti, 2009].

While the sandwich estimate of variance provides asymptotically correct inference, the small sample confidence interval coverage is typically quite poor. This has been attributed to the increased variability in the sandwich estimator over other, less robust variance estimators [Kauermann and Carroll, 2001]. Specifically, Efron [1986] in his discussion of Wu [1986] noticed greater variability in the sandwich estimator than in the standard variance estimator in his simulation results. Breslow [1990] also finds greater variability in the sandwich estimator in his simulations of overdispersed Poisson regression. Firth [1992] and McCullagh [1992] in their review of Liang et al. [1992] both question the small sample efficiency of the sandwich. And Diggle et al. [2009] note that best results are obtained when there are "many experimental units". The result of greater variability in typical small sample scenarios is that confidence intervals may only yield much less than nominal coverage of the pseudo-true parameter. Greater variability and consequently poor confidence interval coverage is the result both of multiple approximations being used in the derivation of confidence intervals based on the sandwich estimate of variance and of multiple quantities being estimated.

Several authors have attempted to adjust the sandwich estimate of variance to improve confidence interval coverage in small sample sizes. The simplest, though not the first, adjustment to the sandwich is to multiply the standard estimator by $\frac{n}{n-p}$ where p is the number of parameters in the model and n is the size of the dataset [Hinkley, 1977]. Following the convention laid out in Hardin [2003], we call this adjustment "hc2". Another earlier attempt to correct for the bias of the sandwich's variance estimate was Horn et al. [1975]. This correction divided the i^{th} diagonal element in the inner, "peanut butter" matrix by $1 - h_{ii}$, where h_{ii} is the i^{th} diagonal element from the hat matrix. The hat matrix refers to the projection matrix generated by the design matrix in linear regression. We call this adjustment "hc1". A jackknife estimator was used by Efron [1982], which is equivalent to dividing the i^{th} diagonal element in the inner, "peanut butter" matrix by $(1-h_{ii})^2$. We call this adjustment "hc3". In Kauermann and Carroll [2000] an adjusted normal quantile

value is computed, instead of adjusting the sandwich directly. We call this adjustment "hc4". Fay and Graubard [2001] compute the degrees of freedom needed for approximately correct $t-$ and $F-$distributions, instead of using asymptotically correct normal and chi-squared distributions. We call this adjustment "hc5".

As remarked upon in Hardin [2003], all of these adjustments to the sandwich variance estimator are "ad hoc solutions". We think a better critique would be that they are "post hoc solutions", since they all attempt to adjust the sandwich estimate of variance after the fact. All of them attempt to inflate confidence interval sizes, but none of them do so in a way that directly addresses the variability inherent in the sandwich estimate of variance. Two of the methods ([Kauermann and Carroll, 2000] and [Fay and Graubard, 2001]) do not focus on adjusting the sandwich estimate of variance. Instead, they modify the standard normal quantiles used in confidence interval construction, or they replace asymptotic normal and chi-squared distributions with Student-t and F distributions. To date, there is no universally-agreed-upon adjustment or alternative to the sandwich variance estimator for small sample sizes that corrects for the its structural problems.

Instead of trying to adjust the sandwich estimate of variance after the fact, we analyze the argument that produces the sandwich. We find three approximating assumptions in this argument. By eliminating one of these approximating assumptions and remembering that the ultimate goal is inference on a parameter, we derive a pivot-based method of inference that produces more accurate confidence intervals in small sample sizes than the sandwich estimate of variance and its "post hoc" adjustments.

The rest of the paper is organized as follows. In Section 2, we derive our pivot-based method by minimizing the number of assumptions used in the standard sandwich-based method. By avoiding plug-in assumptions commonly used with sandwich-based inference, we preserve any mean-variance relationships in the model and improve confidence interval coverage in smaller sample sizes. In Section 3, we give a proof in the one-parameter case that no asymptotic efficiency is lost by using the pivot versus the sandwich. In that same section we follow up with simulation examples that show we can achieve improved coverage in small sample sizes. In Section 4, we extend our proof of asymptotic equivalence to the multiple parameter case, and show a corresponding simulation

exercise using a simple linear regression model. Section 5 takes data from a real-world example, runs a multiple regression on the data, and computes confidence region coverage for all parameters combined.

2 Pivot Method

The argument deriving the sandwich estimate of variance makes three approximating assumptions. In this Section we review those assumptions. Our approach to inference eliminates the plug-in assumption, resulting in better confidence interval coverage at small sample sizes.

We now describe the derivation of the sandwich estimate of variance for a one-parameter model under misspecification. Let $y_1, \ldots, y_n$ represent data values. Let the family of distributions assumed by the analyst be described by the set of densities $\{f(y;\theta) : \theta \in \Theta\}$ with indexing parameter θ, and Θ being a compact subset of R^p. We will write true distribution of the data as density $g(y)$, where $g(y)$ is not necessarily in $\{f(y;\theta) : \theta \in \Theta\}$. Let $\prod_{i=1}^n f(y_i;\theta)$ be the pseudo-likelihood of independent and identically distributed data. Then the pseudo-log-likelihood is given by $\sum_{i=1}^n \log(f(y_i;\theta))$. We denote the derivative of the pseudo-log-likelihood with respect to θ by $\sum_{i=1}^n l_\theta(y_i;\theta)$ where the subscript denotes the derivative. The second derivative of the pseudo-log-likelihood with respect to θ is similarly denoted by $\sum_{i=1}^n l_{\theta\theta}(y_i;\theta)$. We let $\hat{\theta}$ denote the maximum misspecified likelihood estimator. This is simply the estimator we compute when using maximum likelihood methods on a possibly misspecified likelihood. We use θ^* to represent the Kullback-Leibler minimizing parameter, which we refer to as the pseudo-true parameter. Then by a mean value lemma for random variables proven in Jennrich [1969] there exists $\bar{\theta}$, a weighted average of $\hat{\theta}$ and θ^*, that satisfies the following equation

$$-\frac{1}{\sqrt{n}}\sum_{i=1}^{n} l_\theta(y_i,\theta^*) = \frac{1}{\sqrt{n}}\left[\sum_{i=1}^{n} l_{\theta\theta}(y_i,\bar{\theta})\right](\hat{\theta}-\theta^*) \tag{1}$$

From here we solve for $\sqrt{n}(\hat{\theta}-\theta^*)$ and compute the variance of its limiting distribution. In the process, one must make three approximating assumptions:

1. $\frac{1}{\sqrt{n}}\sum_{i=1}^n l_\theta(y_i,\theta^*) \dot{\sim} N(0, B(\theta^*))$

2. $\frac{1}{n}\sum_{i=1}^{n}[l_\theta(y_i,\theta^*)]^2 = B_n(\theta^*) \approx E\{[l_\theta(y,\theta^*)]^2\} = B(\theta^*)$

3. $\hat{\theta} \approx \theta^*$.

Assumption 1 is a distributional assumption. Since we assume a possibly misspecified model, we cannot state the exact distribution of the score function. Instead, we rely on the Central Limit Theorems for an approximate distribution. Assumption 2 is an empirical assumption. Again because we assume a possibly misspecified model we cannot compute the expectation in that line. We use the empirical approximation instead. Assumptions 3 is the plug-in assumption. We do not know the true value of the pseudo-true parameter, so we approximate the value with its estimator.

If we let $A_n(\bar{\theta}) = \frac{1}{n}\sum l_{\theta\theta}(y_i,\bar{\theta})$ then the argument proceeds as follows:

$$\frac{1}{\sqrt{n}}\sum l_\theta(y_i,\theta^*) \dot{\sim} N(0, B(\theta^*)) \tag{2}$$

is directly from assumption (i).

$$A_n(\bar{\theta})\sqrt{n}(\hat{\theta}-\theta^*) \dot{\sim} N(0, B(\theta^*)) \tag{3}$$

is derived by replacing the score function in assumption (i) with the right-hand side of (1).

$$B^{-1/2}(\theta^*)A_n(\bar{\theta})\sqrt{n}(\hat{\theta}-\theta^*) \dot{\sim} N(0,1) \tag{4}$$

is computed by multiplying both sides by $B^{-1/2}(\theta^*)$.

$$B_n^{-1/2}(\theta^*)A_n(\bar{\theta})\sqrt{n}(\hat{\theta}-\theta^*) \dot{\sim} N(0,1) \tag{5}$$

results from applying assumption (ii).

$$B_n^{-1/2}(\hat{\theta})A_n(\hat{\theta})\sqrt{n}(\hat{\theta}-\theta^*) \dot{\sim} N(0,1) \tag{6}$$

uses assumptions (iii) and (iv), replacing θ^* with $\hat{\theta}$ in A_n. Recall that $\bar{\theta}$ is a linear combination of $\hat{\theta}$ and θ^*, so replacing $\bar{\theta}$ with $\hat{\theta}$ in $B_n^{-1/2}$ is equivalent to replacing θ^* with $\hat{\theta}$.

Our goal is to eliminate one or more of these approximating assumptions in order to improve accuracy in confidence intervals about θ^*. In our derivation we also start with line (2), but we remind ourselves that finding the asymptotic distribution of $\sqrt{n}(\hat{\theta} - \theta^*)$ is only a tool for creating confidence intervals and is not our end goal. Our goal is to improve confidence interval coverage about θ^*.

Since we are operating under the assumption of model misspecification, we have no knowledge of the exact distribution of $\frac{1}{\sqrt{n}} \sum_{i=1}^{n} l_\theta(y_i, \theta^*)$. The best we can do is to keep the assumption of approximate normality by appealing to the Central Limit Theorem and using assumption (i) above. This may be a crude approximation, but it is the only distributional result we have. Model misspecification also means we cannot compute the asymptotic variance $B(\theta)$, so we estimate it with its empirical approximation, $B_n(\theta)$, which is assumption (ii) from above. With these assumptions alone, we can still perform inference using an approximate pivotal quantity. The pivot and its asymptotically correct distribution is derived now.

$$\frac{1}{\sqrt{n}} \sum l_\theta(y_i, \theta^*) \dot{\sim} N(0, B(\theta^*)) \tag{7}$$

is the same as line (2) and

$$B^{-1/2}(\theta^*) \frac{1}{\sqrt{n}} \sum l_\theta(y_i, \theta^*) \dot{\sim} N(0, 1) \tag{8}$$

is the result of multiplying both sides by $B^{-1/2}(\theta^*)$. Applying Assumption 2 we get

$$B_n^{-1/2}(\theta^*) \frac{1}{\sqrt{n}} \sum l_\theta(y_i, \theta^*) \dot{\sim} N(0, 1) \tag{9}$$

And the explicit, approximate pivot is

$$t(y_1, \ldots, y_n; \theta^*) = \left\{ \frac{1}{n} \sum_{i=1}^{n} [l_\theta(y_i, \theta^*)]^2 \right\}^{-1/2} \frac{1}{\sqrt{n}} \sum_{i=1}^{n} l_\theta(y_i, \theta^*) \dot{\sim} N(0, 1). \tag{10}$$

With (10), one can perform a linear search on θ^* to find confidence interval boundaries given a confidence level of $100(1 - \alpha)\%$. Find the θ-values that makes the left-hand side of line (10) equal

to the desired standard normal quantiles, and that is the confidence interval. Using this method, we have eliminated the plug-in approximating assumption required to perform inference using the sandwich estimate of variance.

3 One-Parameter Asymptotics and Simulation Example

In this section we will do two things, both in the univariate parameter case. First we show that the pivot loses no efficiency with respect to the sandwich. This is important because if the pivot were only better than the sandwich at small sample sizes, but lost efficiency with respect to the sandwich at larger sample sizes, we would have no way of determining which method to use in order to perform the best inference possible. Second, we give a couple of simulation examples comparing the pivot method with the sandwich and its post hoc adjustments. Our examples show that the pivot method performs favorably.

3.1 One-Parameter Asymptotics

We now show that the confidence interval procedure based on pivot (10) is asymptotically as efficient as the confidence interval based on the plug-in sandwich estimator. In the context of confidence interval evaluation, efficiency is judged based on relative confidence interval width given a particular confidence level. For a one-sided confidence interval, the equivalent comparison is of the values of the finite boundary. We show in this subsection that

$$\sqrt{n}|\theta_{n,1} - \theta_{n,2}| \rightarrow_{a.s.} 0 \tag{11}$$

where $\theta_{n,1}$ and $\theta_{n,2}$ that are defined by the equalities

$$\sqrt{n}B_n^{-1/2}(\theta_{n,1})\frac{1}{n}\sum l_\theta(y_i, \theta_{n,2}) = z_{1-\alpha} \tag{12}$$

$$\sqrt{n}B_n(\hat{\theta})^{-1/2}A_n(\hat{\theta})(\hat{\theta} - \theta_{n,2}) = z_{1-\alpha}, \tag{13}$$

which yield the boundaries of the one-sided confidence interval using the pivot, (12), and the sandwich, (13), respectively. To find a two-sided interval, we would set each equation above to $z_{1-\alpha/2}$ and $z_{\alpha/2}$, resulting in four total equations, and then solve for the interval endpoints.

In order to facilitate computation, we will substitute Equation (1) into Equation (12) to get

$$\sqrt{n}B_n^{-1/2}(\theta_{n,1})\frac{1}{n}\sum l_\theta(y_i,\theta_{n,2}) = \sqrt{n}B_n(\theta_{n,1})^{-1/2}A_n(\bar{\theta})(\hat{\theta}-\theta_{n,1}) = z_{1-\alpha}. \tag{14}$$

The reason for this approach is intuitive. If two estimators have asymptotically normal distributions, then computing the relative efficiency of the estimators is equivalent to comparing confidence interval widths. If the widths are the same, then the estimators are equally efficient. If the widths are different, then the shorter confidence interval is more efficient.

Expressions that we will use for each of $\theta_{n,1}$ and $\theta_{n,2}$ are derived from Equations (13) and (14)

$$\theta_{n,1} = \hat{\theta}_n - \frac{A_n((1-b_n)\hat{\theta}_n + b_n\theta_{1,n})^{-1}B_n(\theta_{n,1})^{1/2}z_{1-\alpha}}{\sqrt{n}} \tag{15}$$

$$\theta_{n,2} = \hat{\theta}_n - \frac{A_n(\hat{\theta})^{-1}B_n(\hat{\theta})^{1/2}z_{1-\alpha}}{\sqrt{n}} \tag{16}$$

where b_n is the value between 0 and 1 that satisfies the Lemma 3 equality in Jennrich [1969]. Notice that we did not solve for $\theta_{n,1}$ explicitly. This is not a problem. We are considering asymptotic behavior, and we prove in the appendices that $\theta_{1,n} \rightarrow_{a.s.} \theta^*$. Now consider the scaled difference

$$\begin{aligned}\sqrt{n}|\theta_{n,1}-\theta_{n,2}| = \sqrt{n}\Bigg|&\hat{\theta}_n - \frac{A_n((1-b_n)\hat{\theta}_n + b_n\theta_{n,1})^{-1}B_n(\theta_{n,1})^{1/2}z_{1-\alpha}}{\sqrt{n}} \\ &- \left(\hat{\theta}_n - \frac{A_n(\hat{\theta})^{-1}B_n(\hat{\theta})^{1/2}z_{1-\alpha}}{\sqrt{n}}\right)\Bigg| \end{aligned} \tag{17}$$

$$\tag{18}$$

where we replace the left-hand side with the equivalent expressions from (15) and (16). With some

algebra we get

$$= \sqrt{n} \left| \frac{A_n(\hat{\theta})^{-1} B_n(\hat{\theta})^{1/2} z_{1-\alpha}}{\sqrt{n}} - \frac{A_n((1-b_n)\hat{\theta}_n + b_n\theta_{n,1})^{-1} B_n(\theta_{n,1})^{1/2} z_{1-\alpha}}{\sqrt{n}} \right| \tag{19}$$

and

$$= \left| A_n(\hat{\theta})^{-1} B_n(\hat{\theta})^{1/2} z_{1-\alpha} - A_n((1-b_n)\hat{\theta}_n + b_n\theta_{1,n})^{-1} B_n(\theta_{n,1})^{1/2} z_{1-\alpha} \right|. \tag{20}$$

$$\tag{21}$$

And finally we have

$$\left| A_n(\hat{\theta})^{-1} B_n(\hat{\theta})^{1/2} z_{1-\alpha} - A_n((1-b_n)\hat{\theta}_n + b_n\theta_{1,n})^{-1} B_n(\theta_{n,1})^{1/2} z_{1-\alpha} \right| \to_{a.s.} \tag{22}$$

$$z_{1-\alpha} \left| A(\theta^*)^{-1} B(\theta^*)^{1/2} - A(\theta^*)^{-1} B(\theta^*)^{1/2} \right| \tag{23}$$

$$= 0. \tag{24}$$

These last lines are the consequence of several convergence results, the first being $\theta_{n,1} \to_{a.s.} \theta^*$, which is shown in the appendices. By the strong law of large numbers, we have $A_n(\hat{\theta}) \to_{a.s.} A(\theta^*)$, $A_n(\bar{\theta}) \to_{a.s.} A(\theta^*)$, $B_n(\hat{\theta}) \to_{a.s.} B(\theta^*)$, $B_n(\theta_{n,1}) \to_{a.s.} B(\theta^*)$, and by Mann-Wald's Theorem we have line (23) above. Since the $\sqrt{n}$-distance between confidence interval boundaries converges to zero, the confidence intervals generated by the pivot and the sandwich are asymptotically equivalent. Hence, the two procedures for generating confidence intervals are asymptotically equally efficient.

3.2 One-Parameter Small Sample Comparisons

In this section we examine two, one-parameter examples comparing confidence interval performance between intervals created using the likelihood pivot and intervals created using the sandwich estimate of variance. The first example assumes a Poisson likelihood for count data. The second example assumes a regression through the origin model with one response variable and one covariate. Our choice of examples is meant to showcase the diverse applications of the likelihood pivot. In our two examples we show that our method works with count data, with continuous data, in

the presence of heteroscedasticity, and in a regression setting. Broad applicability is one of the strengths of the likelihood pivot method for generating confidence intervals.

3.3 Poisson Data

In this example, we assume a working model

$$y_i \overset{i.i.d.}{\sim} Pois(\theta) \qquad i = 1, \ldots, n \tag{25}$$

and we wish to perform inference on the parameter of interest θ. A typical misspecification would be assuming the above model when an overdispersed Poisson is the truth. We can generate such a distribution using the Negative Binomial distribution, since the variance of a Negative Binomial random variable can differ from its mean. The Negative Binomial is what we used as our true data-generating distribution in this example.

In this model we have the following pieces of Equations (1) and (10)

$$l_\theta(\theta^*) = \frac{1}{\theta^*} \sum_{i=1}^{n} y_i - n \tag{26}$$

$$\hat{\theta} = \frac{1}{n} \sum_{i=1}^{n} y_i \tag{27}$$

$$l_{\theta\theta}(\bar{\theta}) = \frac{-1}{(\bar{\theta})^2} \sum_{i=1}^{n} y_i \tag{28}$$

We then plug those into the pivot expression to get

$$\frac{1}{\sqrt{n}} B_n(\theta^*)^{-1/2} l_\theta(y_1, \ldots, y_n; \theta^*) = \frac{1}{\sqrt{n}} \frac{\frac{1}{\theta^*} \sum_{i=1}^{n} y_i - n}{\sqrt{\frac{1}{(\theta^*)^2} \frac{1}{n} \sum_{i=1}^{n} (y_i - \theta^*)^2}} \tag{29}$$

$$= \sqrt{n} \frac{\frac{1}{\theta^*} \frac{1}{n} \sum_{i=1}^{n} y_i - 1}{\frac{1}{|\theta^*|} \sqrt{\frac{1}{n} \sum_{i=1}^{n} (y_i - \theta^*)^2}} \tag{30}$$

$$= \sqrt{n} \frac{\frac{1}{n} \sum_{i=1}^{n} y_i - \theta^*}{\sqrt{\frac{1}{n} \sum_{i=1}^{n} (y_i - \theta^*)^2}} \tag{31}$$

and we perform a linear search on the parameter in this pivot to find our confidence interval.

The sandwich estimate of variance is calculated as

$$A_n(\hat{\theta})^{-1}B_n(\hat{\theta})A_n(\hat{\theta})^{-1} = \left[\frac{-1}{n}\sum_{i=1}^{n}\frac{y_i}{\hat{\theta}^2}\right]^{-1}\frac{1}{n}\sum_{i=1}^{n}\left(\frac{y_i}{\hat{\theta}}-1\right)^2\left[\frac{-1}{n}\sum_{i=1}^{n}\frac{y_i}{\hat{\theta}^2}\right]^{-1} \tag{32}$$

$$= \left[\frac{-1}{n}\sum_{i=1}^{n}\frac{y_i}{\hat{\theta}^2}\right]^{-2}\frac{1}{n}\sum_{i=1}^{n}\left(\frac{y_i}{\hat{\theta}}-1\right)^2 \tag{33}$$

$$= \left[\frac{-\hat{\theta}}{\hat{\theta}^2}\right]^{-2}\frac{1}{n}\sum_{i=1}^{n}\left(\frac{y_i}{\hat{\theta}}-1\right)^2 \tag{34}$$

$$= \left[\frac{-1}{\hat{\theta}}\right]^{-2}\frac{1}{n}\sum_{i=1}^{n}\left(\frac{y_i}{\hat{\theta}}-1\right)^2 \tag{35}$$

$$= \hat{\theta}^2\frac{1}{n}\sum_{i=1}^{n}\left(\frac{y_i}{\hat{\theta}}-1\right)^2 \tag{36}$$

$$= \hat{\theta}^2\frac{1}{n\hat{\theta}^2}\sum_{i=1}^{n}\left(y_i-\hat{\theta}\right)^2 \tag{37}$$

$$= \frac{1}{n}\sum_{i=1}^{n}\left(y_i-\hat{\theta}\right)^2 \tag{38}$$

From the above, we see that the sandwich estimate of variance for $\hat{\theta}$ is $\frac{1}{n}\sum_{i=1}^{n}(y_i-\hat{\theta})^2$, while the corresponding variance factor in the likelihood pivot is $\frac{1}{n}\sum_{i=1}^{n}(y_i-\theta^*)^2$. The only difference between the formulas in this example is the substitution of θ^* in the pivot quantity with $\hat{\theta}$ in the sandwich estimate of variance.

For the simulation study, I used a negative binomial with mean 3 and variance 3.9 as the true data-generating distribution. I then generated $10,000$ samples of size ranging from 10 to 100. For each sample, I constructed asymptotically correct 95% confidence intervals based on the sandwich estimate of variance and its post hoc adjustments. I also constructed a confidence interval by conducting a linear search on the pivot to find the parameter values that caused the pivot to equal the 2.5% and 97.5% quantiles of the standard normal distribution. The results can be seen in Figure 1

The pivot-based confidence intervals perform well at all sample sizes. Coverage hovers around 95% throughout all sample sizes tested. Standard maximum likelihood theory undercovers the true parameter, a reflection of the fact that we are assuming the mean and variance are equal when

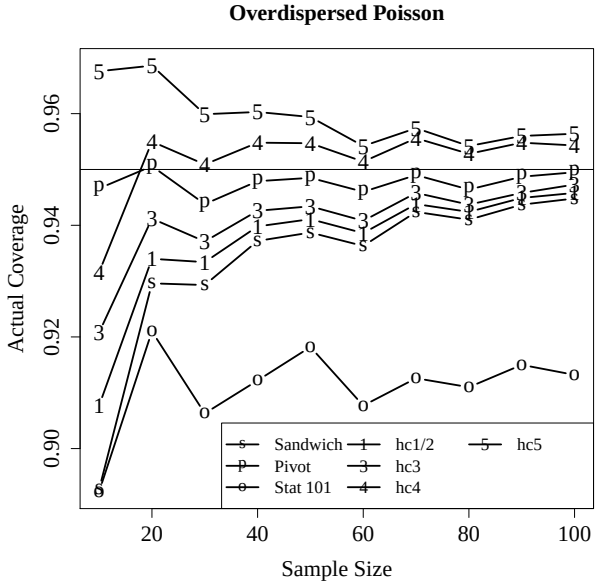

Figure 1: Results from simulation study using sandwich-based confidence intervals and confidence intervals generated by the likelihood pivot when applied to an assumed Poisson data-generating model.

the variance is actually larger. The sandwich estimate of variance performs as one would expect. It undercovers the true mean in all sample sizes tested, but coverage improves as the sample size increases. The post hoc intervals tend to mirror this as well. The post hoc intervals associated with Fay and Graubard [2001] actually tend to overcover the true parameter. Fay and Graubard [2001] use an adjusted sandwich and the Student t-distribution quantiles instead of normally distributed quantiles. This two-pronged approach might be over-correcting the sandwich's undercoverage, resulting in the overcoverage we see in the simulation results.

In this simple scenario, we see that all confidence interval construction methods that account for model misspecification perform well (± 1 percentage point), especially as sample sizes increase. While the likelihood pivot performs the best in this simulation, the improvement is not great. This is a reflection of the fact that there is only one parameter in our working model, and hence the number of quantities being estimated in any method is small.

3.3.1 Regression Through the Origin

Regression models are one of the most fundamental statistical models used by practitioners. Any technique that improves on methods using the sandwich estimate of variance must work in a regression setting. In this example, we assume as our working model a two-parameter regression through the origin

$$y_i \sim \theta x_i + \epsilon_i \tag{39}$$

$$\epsilon_i \overset{i.i.d.}{\sim} N(0, \sigma^2) \qquad i = 1, \ldots, n \tag{40}$$

where the parameter of interest is θ and σ^2 is a nuisance parameter. For this working model, the pivot does not depend on σ^2 so it can be unknown. We then consider the behavior of the sandwich-based and pivot-based confidence intervals in the case of heteroscedastic errors. Heteroscedastic errors are a commonly assumed model misspecification, and one of special interest to economists. Specifically the true model is

$$y_i \sim \theta x_i + \epsilon_i \tag{41}$$

$$\epsilon_i \sim N(0, 1+ \mid x_i \mid) \qquad i = 1, \ldots, n \tag{42}$$

where the errors depend directly on the covariate.

Under this model, the log-likelihood, its derivative with respect to θ, and the square of the derivative are

$$\sum_{i=1}^{n} l(y_i, \theta) = -n \log(\sqrt{2\pi}\sigma) + \frac{-1}{2\sigma^2} \sum_{i=1}^{n} (y_i - \theta x_i)^2 \tag{43}$$

$$\sum_{i=1}^{n} l_\theta(y_i, \theta) = \frac{-1}{2\sigma^2} \sum_{i=1}^{n} -2x_i(y_i - \theta x_i) = \frac{1}{\sigma^2} \sum_{i=1}^{n} x_i(y_i - \theta x_i) \tag{44}$$

$$\sum_{i=1}^{n} l_\theta(y_i, \theta)^2 = \frac{1}{\sigma^4} \sum_{i=1}^{n} x_i(y_i - \theta x_i)^2 \tag{45}$$

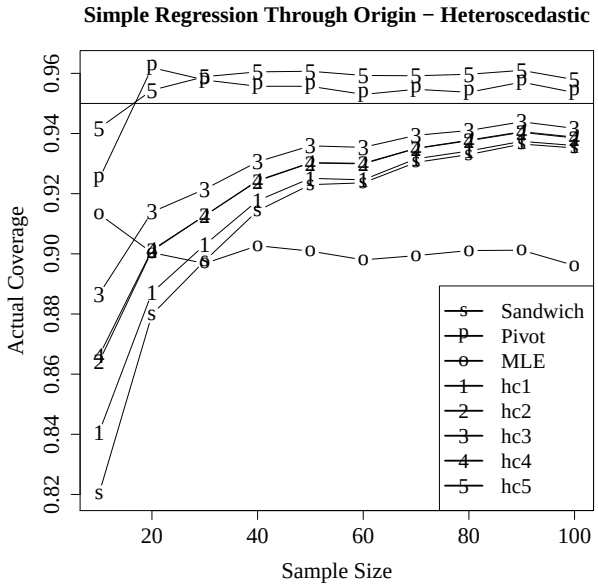

Figure 2: Graph showing actual confidence interval coverage of the true slope parameter when using various methods to compute confidence intervals. For each sample size, 10,000 random datasets were generated, maximum misspecified likelihood estimates were calculated, and confidence intervals were computed.

respectively. Plugging these into line (10), gives us the pivot for this example

$$\left\{\frac{1}{n}\sum_{i=1}^{n}[l_\theta(y_i,\theta^*)]^2\right\}^{-1/2}\frac{1}{\sqrt{n}}\sum_{i=1}^{n}l_\theta(y_i,\theta^*)=\sqrt{n}\frac{\frac{1}{n}\sum_{i=1}^{n}(x_iy_i-\theta^*x_i^2)}{\sqrt{\frac{1}{n}\sum_{i=1}^{n}\left(x_iy_i-\theta^*x_i^2\right)^2}}. \tag{46}$$

This simulation was performed similarly to the previous one. We generated data for samples sizes ranging from 10 to 100. For each sample size, we generated $10,000$ datasets. For each dataset, we computed 95% confidence intervals using standard MLE theory, the sandwich, the pivot, and the post hoc corrections to the sandwich. We then computed whether or not the true value of the parameter was covered by the intervals. The results are displayed graphically in Figure 2.

The standard MLE confidence intervals undercover the pseudo-true parameter throughout the range of sample sizes. Since the assumed model does not account for the variability in errors, we expect confidence intervals to be too small and coverage of the true parameter to be less than nominal. The sandwich-based confidence intervals yield about 82% coverage at the smallest sample size generated, and increase coverage as sample size increases. Since the sandwich has this known

bias, this is not surprising. Three of the five post hoc methods, [Horn et al., 1975], [Hinkley, 1977], and [Efron, 1982], are simply scaled-up versions of the sandwich. As expected they yield coverage similar to, but slightly better than, the unadjusted sandwich. The fourth post hoc method uses the sandwich variance but with a slightly wider quantile spread, e.g. the 98% quantile instead of 97.5% quantile. As such it also covers the true parameter with similar but slightly better coverage than the regular sandwich. The fifth post hoc method yields similar but slightly higher coverage than the pivot. Not only does it use an adjusted sandwich variance estimate, but it also uses Student-t quantiles instead of standard normal. This combination pushes coverage much higher than unaltered sandwich-based coverage. The pivot-based intervals perform the best out of all options tested and yield coverage of the pseudo-true parameter close to 95% throughout the range of sample sizes tested and versus all competing methods.

The results of this subsection match what we saw in Subsection 3.3. Confidence intervals based on the likelihood pivot provide coverage closer to nominal than sandwich-based methods when sample sizes are small. It is no surprise that the post hoc adjustments to the sandwich seem to track one another in Figure 2, since they generally inflate the sandwich in various ways. What is of note is how quickly and consistently the likelihood pivot based intervals reach nominal coverage, especially compared to the other intervals. This has happened in both examples so far.

4 Multi-Parameter Theory and Examples

In this section we consider inference using multi-parameter models. Confidence regions for sets of parameters are important because they yield reasonable combinations of parameters, something that is not readily apparent from multiple, single-parameter confidence intervals. We show how to extend the pivot approach to this situation and compare confidence regions of the pivot approach to those based on the sandwich plug-in estimate of variance. Confidence regions based on the sandwich estimate of variance are values of $\boldsymbol{\theta}$ such that

$$n(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta})^T \left[\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\mathbf{B}_n(\hat{\boldsymbol{\theta}}_n)^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\right](\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}) \leq \chi^2_{p,1-\alpha} \tag{47}$$

where p is the dimension of the parameter of interest and $\chi^2_{p,1-\alpha}$ is the $100(1-\alpha)\%$ quantile of a chi-squared distribution with p degrees of freedom.

Since we are assuming model misspecification, the only distributional result we have is the multivariate central limit theorem. This gives us an approximate distribution for the multivariate pivot

$$\frac{1}{\sqrt{n}}\sum_{i=1}^{n} l_{\boldsymbol{\theta}}(y_i,\boldsymbol{\theta}) \dot{\sim} N(\mathbf{0}_p, \mathbf{B}(\boldsymbol{\theta})). \tag{48}$$

We know from standard theory for multivariate distributions that if

$$\mathbf{Z} \sim N(\mathbf{0}_p, \mathbf{B}) \tag{49}$$

then

$$\mathbf{Z}^T\mathbf{B}^{-1}\mathbf{Z} \sim \chi^2_p. \tag{50}$$

So combining (48) and (50) we have that

$$\left[\frac{1}{\sqrt{n}}\sum_{i=1}^{n} l_{\boldsymbol{\theta}}(y_i,\boldsymbol{\theta})\right]^T \mathbf{B}(\boldsymbol{\theta})^{-1}\left[\frac{1}{\sqrt{n}}\sum_{i=1}^{n} l_{\boldsymbol{\theta}}(y_i,\boldsymbol{\theta})\right] \dot{\sim} \chi^2_p. \tag{51}$$

We can now generate a confidence region through the pivot by finding solutions, $\boldsymbol{\theta}$ to the inequality

$$n\left[\frac{1}{\sqrt{n}}\sum_{i=1}^{n} l_{\boldsymbol{\theta}}(y_i,\boldsymbol{\theta})\right]^T \mathbf{B}(\boldsymbol{\theta})^{-1}\left[\frac{1}{\sqrt{n}}\sum_{i=1}^{n} l_{\boldsymbol{\theta}}(y_i,\boldsymbol{\theta})\right] \leq \chi^2_{p,1-\alpha} \tag{52}$$

We wish to use an expression that more directly comparable to line 47. We will use a multivariate equivalent to Jennrich's Lemma 6 (Mean Value Theorem for functions of random variables) [Jennrich, 1969], to equate $\frac{1}{\sqrt{n}}\sum_{i=1}^{n} l_{\boldsymbol{\theta}}(y_i,\boldsymbol{\theta})$ and $\overline{\mathbf{A}_n(\boldsymbol{\theta})}(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta})$, up to a minus sign. For our analytical work we will now use the equivalent statement

$$n(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta})^T\overline{\mathbf{A}_n(\boldsymbol{\theta})}\mathbf{B}_n(\boldsymbol{\theta})^{-1}\overline{\mathbf{A}_n(\boldsymbol{\theta})}(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}) \leq \chi^2_{p,1-\alpha}. \tag{53}$$

The first difference between the multivariate pivot quantity and the similar sandwich-based quantity is in using $\boldsymbol{\theta}$ instead of $\hat{\boldsymbol{\theta}}_n$ in $\mathbf{B}_n(\cdot)^{-1}$. This is equivalent to eliminating the plug-in assumption from the univariate case. The second is the introduction of $\overline{\mathbf{A}_n(\boldsymbol{\theta})}$. A few words need to be said about this second substitution. Primarily, there is no single point $\boldsymbol{\theta}$ that satisfies Jennrich's Lemma 6 in the multivariate setting. The reason is that there is no mean value theorem for vector-valued functions. $\overline{\mathbf{A}_n(\boldsymbol{\theta})}$ is really a convex combination of matrices as described in Furi and Martelli [1991]. That is, there are $\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_k$ and $\omega_1, \ldots, \omega_k \geq 0$ with $\sum_{i=1}^{k} \omega_i = 1$ such that $\overline{\mathbf{A}_n(\boldsymbol{\theta})} = \sum_{i=1}^{k} \omega_i \mathbf{A}_n(\boldsymbol{\theta}_i)$. Each $\boldsymbol{\theta}_i$ is a point in p-space where each coordinate is taken from either $\hat{\boldsymbol{\theta}}_n$ or $\boldsymbol{\theta}$, thus, $k \leq 2^p$. When we talk about points that satisfy the pivotal quantity, we are referring to solutions $\boldsymbol{\theta}$ and our k points are derived from $\hat{\boldsymbol{\theta}}$ and $\boldsymbol{\theta}$.

4.1 Multi-Parameter Asymptotics

Lehmann [1999] defines multivariate efficiency in terms of the difference of the asymptotic covariance matrices of two estimators. Since we are not computing an explicit estimator, we do not have a covariance matrix, and we must instead consider what it means to be asymptotically equally efficient. In the univariate case, we considered the finite boundary of two one-sided confidence intervals, and equivalently, the length of two-sided intervals. The multivariate equivalent to univariate length is volume. Thus for a given confidence level, two confidence-region-generating procedures are asymptotically equally efficient if their volumes are asymptotically equivalent. In this subsection we will make the previous statement precise, explain why it is reasonable, and state a result in the multiple parameter case that is similar to the result for the one parameter case.

Let us define for a given α the sandwich equality and pivot equality respectively as the or-equal-to option in Inequalities (47) and (52).

Theorem 1. *For any solution, $\breve{\boldsymbol{\theta}}_{n,PV}$, to the pivot equality and a corresponding solution, $\breve{\boldsymbol{\theta}}_{n,SW}$, to the sandwich equality in the same direction from $\hat{\boldsymbol{\theta}}_n$ as $\breve{\boldsymbol{\theta}}_{n,PV}$, $\sqrt{n}\|\breve{\boldsymbol{\theta}}_{n,PV} - \breve{\boldsymbol{\theta}}_{n,SW}\| \rightarrow_{a.s.} 0$ as $n \rightarrow \infty$.*

Proof. In the Appendix. □

In one sense, this result is the multivariate equivalent of the the result in line (11). In the univariate case, the boundary of a one-sided confidence interval is a point. Subsection 3.1 proves that the boundary points for confidence intervals generated by the sandwich estimate of variance and the pivot converge in norm, even when multiplied by $\sqrt{n}$. Here we state that the boundary points of the multivariate confidence region converge pointwise, when pointwise is taken to mean a fixed direction from the sequence of maximum misspecified likelihood estimators $\{\hat{\boldsymbol{\theta}}_n\}$. We also state that this convergence persists under multiplication by $\sqrt{n}$.

In another sense, this result is one way to formalize the statement "the volumes of the confidence regions are asymptotically equivalent". While not computing volumes explicitly, we are demonstrating that the boundaries of the respective confidence regions are converging to each other. If the boundaries of the regions are converging to each other, then the respective volumes must also be converging to each other, which is what it means for the volumes to be asymptotically equivalent. If the volumes are asymptotically equivalent at equal confidence levels, then the confidence region procedures are asymptotically equally efficient.

4.2 Multiparameter Simulation Example

Having compared the sandwich-based pivot and the likelihood pivot asymptotically, we now consider finite sample comparisons. In this example we use as our working model a simple linear regression

$$y_i = \theta_0 + \theta_1 x_i + \epsilon_i \tag{54}$$

$$\epsilon_i \overset{i.i.d.}{\sim} N(0, \sigma^2) \tag{55}$$

where the parameter of interest is $\boldsymbol{\theta} = (\theta_0, \theta_1)^T$ and σ^2 is a nuisance parameter. As in the regression through the origin example, we can assume σ^2 known, but since it cancels in the pivot we could also assume it is an unknown without loss of generality.

We evaluate pivot confidence regions based on the model described by (54) and (55) when the

true data-generating mechanism is

$$y_i = \theta_0 + \theta_1 x_i + \epsilon_i \tag{56}$$

$$\epsilon_i \overset{indep}{\sim} N(0, 1 + |x_i|) \tag{57}$$

with heteroscedastic errors depending on the covariate. Under this model the relevant term of the pseudo-log-likelihood and necessary partial derivatives are given below:

$$\sum_{i=1}^{n} l(y_i, \boldsymbol{\theta}) = -n \log(\sqrt{2\pi}\sigma) + \frac{-1}{2\sigma^2} \sum_{i=1}^{n} (y_i - \theta_0 - \theta_1 x_i)^2 \tag{58}$$

$$\sum_{i=1}^{n} l_{\theta_0}(y_i, \boldsymbol{\theta}) = \frac{1}{2\sigma^2} \sum_{i=1}^{n} (y_i - \theta_0 - \theta_1 x_i) \tag{59}$$

$$\sum_{i=1}^{n} l_{\theta_1}(y_i, \boldsymbol{\theta}) = \frac{1}{2\sigma^2} \sum_{i=1}^{n} x_i (y_i - \theta_0 - \theta_1 x_i). \tag{60}$$

We can write the factors in the pivot inequality (52) as

$$l_{\boldsymbol{\theta}} = \begin{pmatrix} \frac{1}{2n\sigma^2} \sum_{i=1}^{n} (y_i - \theta_0 - \theta_1 x_i) \\ \frac{1}{2n\sigma^2} \sum_{i=1}^{n} x_i (y_i - \theta_0 - \theta_1 x_i) \end{pmatrix} \tag{61}$$

$$\mathbf{B}_n(y, \boldsymbol{\theta}) = \begin{pmatrix} \frac{1}{4n\sigma^4} \sum_{i=1}^{n} (y_i - \theta_0 - \theta_1 x_i)^2 & \frac{1}{4n\sigma^4} \sum_{i=1}^{n} x_i (y_i - \theta_0 - \theta_1 x_i)^2 \\ \frac{1}{4n\sigma^4} \sum_{i=1}^{n} x_i (y_i - \theta_0 - \theta_1 x_i)^2 & \frac{1}{4n\sigma^4} \sum_{i=1}^{n} x_i^2 (y_i - \theta_0 - \theta_1 x_i)^2 \end{pmatrix} \tag{62}$$

and then the pivot is exactly the one as in inequality (52).

A simulation example was performed similarly to the one-parameter examples. We generated data for samples sizes ranging from 10 to 100. For each sample size, we generated $10,000$ datasets. For each dataset, we computed whether or not the true value of the parameter was covered by the specified confidence regions. The results are shown in Figure 3. We only show results for the sandwich, the pivot, and the hc1, hc2, and hc3 sandwich adjustments. The hc4 and hc5 methods were designed for scalar, linear combinations of parameters and are not adaptable to a multivariate setting.

The results in this multi-parameter simulation mirror the results from the one-parameter simula-

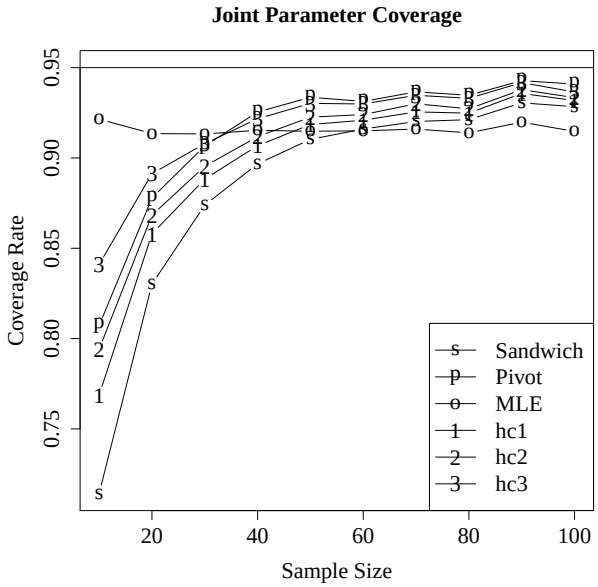

Figure 3: Graph showing confidence interval coverage of the true parameter vector using various methods.

tions. The confidence regions from standard maximum likelihood theory undercover the parameter vector consistently throughout all sample sizes tested. The sandwich-based confidence regions undercover the pseudo-true parameter, with coverage improving as sample size increases. The post hoc methods all provide slightly better coverage than the sandwich, as they all increase the sandwich variance estimate. The pivot-based confidence regions provide comparable coverage to the post hoc methods, providing the best coverage in this simulation in all but the very smallest sample sizes.

We note that all methods which account for misspecification show poor coverage at the smallest sample sizes. This coverage is much worse than what was seen in the one-parameter examples. We think that this is related to the fact that as the dimension of the parameter increases, the number of quantities estimated in the covariance matrix increases with the square of said dimension. As will be seen in the next section, coverage can be quite poor when the parameter vector is even modest in dimension.

5 Data Example

In this section we look at an example using a real dataset as our population. A more realistic comparison of confidence interval methods can be made by using real data as opposed to simulated data. Real world model misspecifications are not just limited to heteroscedastic errors either, so this allows us to compare the methods in the presence of more general model misspecification.

The data comes from the National Health and Nutrition Examination Survey (NHANES) dataset, 2011-2012. We consider a scenario in which a researcher is interested in modeling body-mass index (BMI) as a function of several explanatory variables using a linear regression model. The explanatory variables are taken from variables related to food consumption and activity level, factors that might reasonably affect an individual's BMI.

We use as our population the $7,804$ complete cases in the dataset. The pseudo-true parameters of the model will be the ordinary least-squares estimates applied to these $7,804$ cases. We ran a linear regression of BMI on all the selected explanatory variables to get our pseudo-true parameter values. The model we used was a standard regression model

$$\text{BMI} = \theta_0 + \theta_1 \times \text{kcal} + \theta_2 \times \text{sugar} + \theta_3 \times \text{fat} + \theta_4 \times \text{inact} + \theta_5 \times \text{gender} + \epsilon \tag{63}$$

$$\epsilon \sim N(0, \sigma^2) \tag{64}$$

with "kcal" being the kilocalorie consumption, "inact" being a numerical measurement of inactivity, and ϵ_i representing additional across-subject variability. Next we sampled small datasets from the population to mimic performing multiple small-scale studies. We then ran regressions using those samples, and computed confidence region coverage for all covariate coefficients together. As in the simple linear regression model, we compared the coverage rates of standard maximum likelihood, the sandwich, the pivot, and the hc1, hc2, and hc3 adjustments. The pivot used was identical in structure to the one in the multivariate simulation example, the only difference being the number of covariates was greater in this example. The results can be see in Figure 4.

The standard maximum likelihood method undercovers the pseudo-true vector parameter, showing just under 93% coverage at the higher sample sizes. The sandwich does particularly poorly at

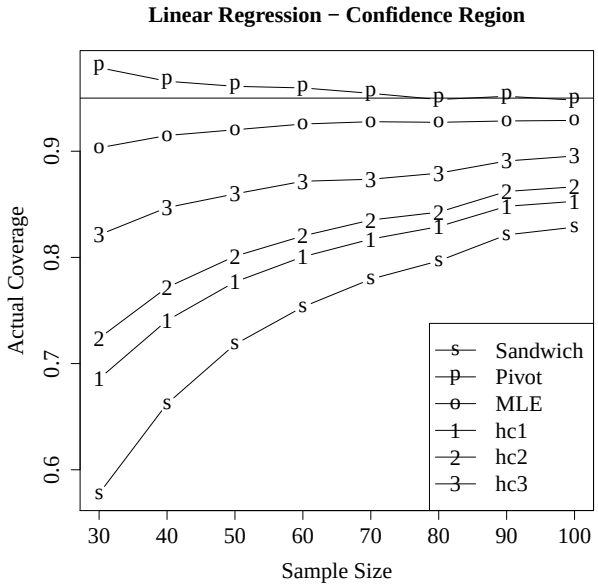

Figure 4: Graph showing confidence interval coverage of the pseudo-true parameter vector value in the NHANES data using various methods. The horizontal line is the 95% line.

first, but improves as sample size increases. The post hoc methods known as hc1, hc2, and hc3 do slightly better than the sandwich, again because they start with the sandwich and increase the variance estimates by varying amounts. The pivot starts out with slightly high coverage, but quickly drops to 95%, and is the only method that achieves 95% coverage within this simulation.

The results of this example are consistent with the results in previous simulations. What this example highlights is that the small sample, model misspecification scenarios under which the likelihood pivot based confidence intervals perform better than sandwich based intervals need not be limited to the examples examined in previous sections. It also highlights the theme running throughout all of these examples: replacing θ^* with $\hat{\theta}$ can seriously lower confidence interval and confidence region coverage. This can cause asymptotically correct inference to differ significantly from nominally correct inference.

6 Discussion

We have presented a pivot-based method of inference that outperforms the sandwich and its variants in small sample sizes. We further showed that this method incurs no efficiency loss at large sample sizes, when compared to sandwich-based inference. Several simulation examples were presented which strengthen our argument that making fewer assumptions leads to better inference.

A natural extension of this work would be to use pivots in a Bayesian context. The seminal paper on Bayesian work with pivots is Monahan and Boos [1992]. The main result of that paper is that likelihoods of pivotal quantities are, in general, not true likelihoods. Consequently, proper Bayesian analysis using pivots is impossible in most cases. However, we believe that some kind of pseudo-Bayesian analysis might be feasible. Our reasoning is simple: for a given set of data, a pivotal quantity should contain some information about the parameter of interest. That information is the basis of the frequentist analysis performed in this paper. A well-designed pseudo-Bayesian analysis might be able to take advantage of that information as well.

A Univariate Results

We will show that $B(\theta^*)$ and $A(\theta^*)$ are positive on Θ. We also need to show that $A_n(\theta^*)$ is positive on this same subset, but due to the uniform central limit theorems, it suffices to show that $A(\theta^*)$ is positive on the required subset. Hence, for large enough n, $A_n(\theta^*)$ is also positive on that subset.

Lemma 1. *Under the assumptions of [White, 1982],* $B(\theta^*) > 0$.

Proof. By definition,

$$B(\theta) = E[l_\theta(y;\theta)^2] \tag{65}$$

where the expectation is taken with respect to the true density. Since we're taking the expectation of the square of a function, we have that

$$B(\theta) \geq 0 \tag{66}$$

for all possible $\theta \in \Theta$.

Assumption A6 in [White, 1982] states that $B(\theta^*)$ is nonsingular. For the one-dimensional case, that means that $B(\theta^*) \neq 0$. Combined with (66), we have $B(\theta^*) > 0$. □

Lemma 2. *Under the assumptions of [White, 1982],* $A(\theta^*) > 0$.

Proof. By definition,

$$A(\theta) = -E[l_{\theta\theta}(y;\theta)]. \tag{67}$$

The case when this is positive is the same as when the expectation of $l_{\theta\theta}$ is negative. Suppose the contrary, namely that θ^* is such that $A(\theta^*) < 0$, or the expectation of $l_{\theta\theta}$ is positive. Since $A_n(\hat{\theta}) \to_{a.s.} A(\theta^*)$, for n large enough, we have that $A_n(\hat{\theta}) < 0$. In other words, the second derivative of the log-likelihood is positive, or in other words that the derivative of the score equation evaluated at θ^* is positive. But this would mean that $\hat{\theta}$ is a minimum of the likelihood equation, and not a maximum. Hence, we restrict ourselves to the cases where $A(\theta^*) \geq 0$.

Assumption A6 in [White, 1982] further states that θ^* is a "regular" point of $A(\theta)$. White [1982] then defines a regular point as one in which $A(\theta)$ has constant rank in an open neighborhood of that point. For the one-dimensional case, we have that if $A(\theta) = 0$, then $A(\theta)$ has rank 0. If $A(\theta) \neq 0$, then $A(\theta)$ has rank 1. So, if $A(\theta^*) \geq 0$, then for the constant rank assumption to hold we must have $A(\theta^*) > 0$ or $A(\theta^*)$ is identically 0 in a neighborhood of θ_*. But if the latter condition holds, then $E[l_\theta]$ is constant in a neighborhood of θ^*, and it no longer has a unique minimum at θ^*. Hence, $A(\theta^*) > 0$. □

We now turn our attention to $\theta_{1,n}$ and $\theta_{2,n}$. We will show that $\theta_{1,n} \to_{a.s.} \theta^*$ and $\theta_{2,n} \to_{a.s.} \theta^*$.

Lemma 3. *Under the assumptions of [White, 1982],* $\theta_{2,n} \to_{a.s.} \theta^*$.

Proof. We can solve directly for $\theta_{2,n}$ in terms of other quantities:

$$\theta_{2,n} = \hat{\theta}_n - \frac{A_n(\hat{\theta}_n)^{-1} B_n(\hat{\theta}_n)^{1/2} z_{1-\alpha}}{\sqrt{n}}. \tag{68}$$

We know from [White, 1982] that $\hat{\theta}_n \to_{a.s.} \theta^*$. By Lemma 4 of [Amemiya, 1973], this means that $B_n(\hat{\theta}_n) \to_{a.s.} B(\theta^*)$ and $A_n(\hat{\theta}_n) \to_{a.s.} A(\theta^*)$. By Mann-Wald this implies that $B_n(\hat{\theta}_n)^{1/2} \to_{a.s.} B(\theta^*)^{1/2}$ and $A_n(\hat{\theta}_n)^{-1} \to_{a.s.} A(\theta^*)^{-1}$. $z_{1-\alpha}$ is a constant. And $\sqrt{n}$ grows arbitrarily large with increasing n. Thus, we have $\theta_{2,n} \to_{a.s.} \theta^* - 0 = \theta^*$. □

Lemma 4. *For all $\epsilon > 0$ there exists N, A_{max}, A_{min}, B_{max}, and B_{min} such that for all $n > N$,*

$$A_{min} - \epsilon \leq A_n(\bar{\theta}_n)^{-1} \leq A_{max} + \epsilon \tag{69}$$

$$B_{min} - \epsilon \leq B_n(\theta_{1,n})^{1/2} \leq B_{max} + \epsilon. \tag{70}$$

Proof. At this point, we remind the reader of some of the standard regularity assumptions generally made in this subject. Specifically, we will assume that the θ-space, Θ, is a compact subset of Euclidean space. We will also rely on Assumption A5 in [White, 1982] to ensure the continuity of $B(\theta)$ and $A(\theta)$. Furthermore, we will be able to apply a uniform law of large numbers to $B_n(\theta)$ and $A_n(\theta)$. A full list of standard regularity assumptions can be found in [White, 1982].

Additionally, we need to be working in a convex set. Technically, Θ is only a compact set in Euclidean space. But, being closed and bounded, we can enclose Θ within a possibly larger, convex, still compact set Θ_{convex}. We will work within Θ_{convex} to set our bounds. We also need to assume that $B_n(\theta)^{1/2}$, $B(\theta)^{1/2}$, $A_n(\theta)^{-1}$ and $A(\theta)^{-1}$ are continuous functions on Θ_{convex}.

Now, since $B_n(\theta)^{1/2}$, $B(\theta)^{1/2}$, $A_n(\theta)^{-1}$ and $A(\theta)^{-1}$ are continuous functions on $\theta \in \Theta_{convex}$ and Θ_{convex} is compact, those functions each map to compact subsets of $\mathbb{R}$, thus they attain their maximum and minimum on the image sets. That is, there are values $\theta_{min,B^{1/2}}$ and $\theta_{max,B^{1/2}}$ such that $0 < B(\theta_{min,B^{1/2}})^{1/2} \leq B(\theta)^{1/2} \leq B(\theta_{max,B^{1/2}})^{1/2} < \infty$, $\forall \theta \in \Theta_{convex}$. We can define $\theta_{min,A^{-1}}$, $\theta_{max,A^{-1}}$, $\theta_{min,A_n^{-1}}$, and $\theta_{max,A_n^{-1}}$, $\theta_{min,B_n^{1/2}}$, and $\theta_{max,B_n^{1/2}}$ similarly.

We will proceed to show the details for the bound on $A_n(\bar{\theta}_n)^{-1}$ knowing that the argument is similar for $B_n(\theta_{1,n})^{1/2}$. From the convexity of Θ_{convex}, we know that $\bar{\theta}_n = (1-b)\hat{\theta}_n + b\theta_{1,n} \in \Theta_{convex}$ for some $b \in [0,1]$, and hence $A_n((1-b)\hat{\theta}_n + b\theta_{1,n})^{-1} \leq A_n(\theta_{max,A_n^{-1}})^{-1}$. Since we can apply the

uniform law of large numbers to $A_n(\theta)^{-1}$, we have that

$$\sup_{\theta \in \Theta_{convex}} |A_n(\theta)^{-1} - A(\theta)^{-1}| \rightarrow_{a.s.} 0. \tag{71}$$

What this means is that with probability one, $\forall \epsilon > 0$, $\exists N_\epsilon$ such that $\forall n \geq N_\epsilon$, $\forall \theta \in \Theta_{convex}$, $|A_n(\theta) - A(\theta)| < \epsilon$. So, if we fix an $\epsilon > 0$, there is some N_ϵ such that for all $n \geq N_\epsilon$ we have

$$A_n((1-b)\hat{\theta}_n + b\theta_{1,n})^{-1} \leq A_n(\theta_{max,A_n^{-1}})^{-1} \tag{72}$$

$$\leq A(\theta_{max,A_n^{-1}})^{-1} + \epsilon \tag{73}$$

$$\leq A(\theta_{max,A^{-1}})^{-1} + \epsilon. \tag{74}$$

Line (72) comes from the definition of $\theta_{max,A_n^{-1}}$. Line (73) is from the uniform law of large numbers. And line (74) is from the definition of $\theta_{max,A^{-1}}$.

Now, we can say that for a fixed $\epsilon > 0$, there exists an $N_{\epsilon,B^{1/2}}$ for $B_n(\theta)^{1/2}$ and an $N_{\epsilon,A_n^{-1}}$ for $A(\theta)^{-1}$ such that for $n \geq max(N_{\epsilon,B^{1/2}}, N_{\epsilon,A_n^{-1}})$

$$A_n((1-b)\hat{\theta}_n + b\theta_{1,n})^{-1} \leq A(\theta_{max,A^{-1}})^{-1} + \epsilon \tag{75}$$

$$B_n(\theta_{1,n})^{1/2} \leq B(\theta_{max,B^{1/2}})^{1/2} + \epsilon. \tag{76}$$

By the same style of argument for n sufficiently large, we have additionally that

$$A_n((1-b)\hat{\theta}_n + b\theta_{1,n})^{-1} \geq A(\theta_{min,A^{-1}})^{-1} - \epsilon \tag{77}$$

$$B_n(\theta_{1,n})^{1/2} \geq B(\theta_{min,B^{1/2}})^{1/2} - \epsilon. \tag{78}$$

If we let $A_{max} = A(\theta_{max,A^{-1}})^{-1}$, $B_{max} = B(\theta_{max,B^{1/2}})^{1/2}$, $A_{min} = A(\theta_{min,A^{-1}})^{-1}$, $B_{min} = B(\theta_{min,B^{1/2}})^{1/2}$ then the result is proved. □

Lemma 5. $\theta_{1,n} \rightarrow_{a.s.} \theta^*$

Proof. $\theta_{1,n}$ solves the following equation

$$\sqrt{n}B_n(\theta_{1,n})^{-1/2}\frac{1}{n}\sum_{i=1}^{n} l_\theta(y_i;\theta_{1,n}) = z_{1-\alpha}. \tag{79}$$

Going back to Jennrich's Lemma, we see that the result still holds replacing θ^* with $\theta_{1,n}$. Thus, we can investigate

$$\sqrt{n}B_n(\theta_{1,n})^{-1/2}A_n(\bar{\theta})(\hat{\theta}-\theta_{1,n}) = z_{1-\alpha} \tag{80}$$

instead.

Since $\bar{\theta}$ lies on the line segment containing $\hat{\theta}$ and $\theta_{1,n}$, we can write $\bar{\theta} = (1-b)\hat{\theta} + b\theta_{1,n}$ for some $b \in [0,1]$. This means our equation can be re-written as

$$\sqrt{n}B_n(\theta_{1,n})^{-1/2}A_n((1-b)\hat{\theta}+b\theta_{1,n})(\hat{\theta}-\theta_{1,n}) = z_{1-\alpha}, \qquad b \in [0,1]. \tag{81}$$

Clearly, we cannot solve this equation directly for $\theta_{1,n}$ in the general case. Instead we perform a linear search to find the value of $\theta_{1,n}$ that satisfies Equation (81).

We do some minor algebra to the previous equation

$$\hat{\theta}_n - \theta_{1,n} = \frac{z_{1-\alpha}}{\sqrt{n}}B_n(\theta_{1,n})^{1/2}A_n((1-b)\hat{\theta}_n + b\theta_{1,n})^{-1} \tag{82}$$

This means that for n sufficiently large, we have

$$-\infty < \frac{z_{1-\alpha}}{\sqrt{n}}(B(\theta_{min,B^{1/2}})^{1/2}-\epsilon)(A(\theta_{min,A^{-1}})^{-1}-\epsilon) \tag{83}$$

$$\leq \hat{\theta}_n - \theta_{1,n} \leq \frac{z_{1-\alpha}}{\sqrt{n}}(B(\theta_{max,B^{1/2}})^{1/2}+\epsilon)(A(\theta_{max,A^{-1}})^{-1}+\epsilon) < \infty. \tag{84}$$

But the products $z_{1-\alpha}(B(\theta_{min,B^{1/2}})^{1/2}-\epsilon)(A(\theta_{min,A^{-1}})^{-1}-\epsilon)$, and $z_{1-\alpha}(B(\theta_{max,B^{1/2}})^{1/2}+\epsilon)(A(\theta_{max,A^{-1}})^{-1}+\epsilon)$ are constants, so the middle term converges to zero as $n \to \infty$ by the Squeeze Theorem. And since $\hat{\theta}_n \to_{a.s.} \theta^*$, we have that $\theta_{1,n} \to_{a.s.} \theta^*$ as well. □

B Multivariate Results

Lemma 6. *Suppose* $\mathbf{A}_1, \ldots, \mathbf{A}_k, \mathbf{A}$ *are matrices such that* $\|\mathbf{A}_i - \mathbf{A}\| < \epsilon$ *for* $i = 1, \ldots, k$ *and for some* $\epsilon > 0$. *Let* $\omega_1, \ldots, \omega_k$ *be weights such that* $\omega_i \geq 0$ *for* $i = 1, \ldots, k$ *and* $\sum_{i=1}^k \omega_i = 1$. *Then* $\|\sum_{i=1}^k \omega_i \mathbf{A}_i - \mathbf{A}\| < \epsilon$.

Proof. We have

$$\|\sum_{i=1}^k \omega_i \mathbf{A}_i - \mathbf{A}\| = \|\sum_{i=1}^k \omega_i (\mathbf{A}_i - \mathbf{A})\| \tag{85}$$

$$\leq \sum_{i=1}^k \|\omega_i (\mathbf{A}_i - \mathbf{A})\| \tag{86}$$

$$= \sum_{i=1}^k \omega_i \|(\mathbf{A}_i - \mathbf{A})\| \tag{87}$$

$$\leq \sum_{i=1}^k \omega_i \epsilon \tag{88}$$

$$= \epsilon \tag{89}$$

□

Lemma 7. *For every* $\epsilon > 0$ *there is a ball of radius* δ *around* $\boldsymbol{\theta}^*$, *called* $Ball_\delta(\boldsymbol{\theta}^*)$, *such that if* $\boldsymbol{\theta}_0, \boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_k \in Ball_\delta(\boldsymbol{\theta}^*)$ *and* $\mathbf{A}(\bar{\boldsymbol{\theta}}) = \sum_{i=1}^k \omega_i \mathbf{A}(\boldsymbol{\theta}_i)$ *a convex combination, then* $\|\mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}^*)\mathbf{A}(\boldsymbol{\theta}^*) - \mathbf{A}(\bar{\boldsymbol{\theta}})\mathbf{B}(\boldsymbol{\theta}_0)\mathbf{A}(\bar{\boldsymbol{\theta}})\| < \epsilon$.

Proof. First, we note the constants $A^* = \|\mathbf{A}(\boldsymbol{\theta}^*)\|$ and $B^* = \|\mathbf{B}(\boldsymbol{\theta}^*)\|$.

Next, we use the continuity of $\mathbf{B}(\boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$ to see that for all $\epsilon > 0$, there is a $\delta_B > 0$ such that if $\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\| < \delta_B$ then $\|\mathbf{B}(\boldsymbol{\theta}) - \mathbf{B}(\boldsymbol{\theta}^*)\| < \epsilon/(3A^*)$.

We now set $B_{bnd} = \sup_{\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\| < \delta_B} \|\mathbf{B}(\boldsymbol{\theta})\|$ and similarly set $A_{bnd} = \sup_{\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\| < \delta_B} \|\mathbf{A}(\boldsymbol{\theta})\|$.

This implies that for all $\epsilon > 0$, there is a $\delta_{B'} > 0$ such that if $\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\| < \delta_{B'}$ then $\|\mathbf{B}(\boldsymbol{\theta}) - \mathbf{B}(\boldsymbol{\theta}^*)\| < \epsilon/(3A^* A_{bnd})$

Now use the continuity of $\mathbf{A}(\boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$ to see that for all $\epsilon > 0$, there is a $\delta_A > 0$ such that if $\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\| < \delta_A$ then $\|\mathbf{A}(\boldsymbol{\theta}) - \mathbf{A}(\boldsymbol{\theta}^*)\| < \epsilon/(3 * \max(A^* \cdot B^*, A_{bnd} \cdot B_{bnd}))$.

Pick $\epsilon > 0$, and let $\delta = \min(\delta_A, \delta_B, \delta_{B'})$. This means that $B_{bnd} \geq \sup_{\|\boldsymbol{\theta}-\boldsymbol{\theta}^*\|<\delta} \|\mathbf{B}(\boldsymbol{\theta})\|$ and likewise $A_{bnd} \geq \sup_{\|\boldsymbol{\theta}-\boldsymbol{\theta}^*\|<\delta} \|\mathbf{A}(\boldsymbol{\theta})\|$.

Consider $Ball_\delta(\boldsymbol{\theta}^*) = \{\boldsymbol{\theta} \mid \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\| < \delta\}$. Now select any $\boldsymbol{\theta}_0, \boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_k \in Ball_\delta(\boldsymbol{\theta}^*)$. And let $\mathbf{A}(\bar{\boldsymbol{\theta}}) = \sum_{i=1}^k \omega_i \mathbf{A}(\boldsymbol{\theta}_i)$ a convex combination. Then we have

$$\|\mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}^*)\mathbf{A}(\boldsymbol{\theta}^*) - \mathbf{A}(\bar{\boldsymbol{\theta}})\mathbf{B}(\boldsymbol{\theta}_0)\mathbf{A}(\bar{\boldsymbol{\theta}})\| \tag{90}$$

$$= \|\mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}^*)\mathbf{A}(\boldsymbol{\theta}^*) - \mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}^*)\mathbf{A}(\bar{\boldsymbol{\theta}}) + \mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}^*)\mathbf{A}(\bar{\boldsymbol{\theta}}) - \mathbf{A}(\bar{\boldsymbol{\theta}})\mathbf{B}(\boldsymbol{\theta}_0)\mathbf{A}(\bar{\boldsymbol{\theta}})\| \tag{91}$$

$$\leq \|\mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}^*)\mathbf{A}(\boldsymbol{\theta}^*) - \mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}^*)\mathbf{A}(\bar{\boldsymbol{\theta}})\| + \|\mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}^*)\mathbf{A}(\bar{\boldsymbol{\theta}}) - \mathbf{A}(\bar{\boldsymbol{\theta}})\mathbf{B}(\boldsymbol{\theta}_0)\mathbf{A}(\bar{\boldsymbol{\theta}})\| \tag{92}$$

$$\leq \|\mathbf{A}(\boldsymbol{\theta}^*)\|\|\mathbf{B}(\boldsymbol{\theta}^*)\|\|\mathbf{A}(\boldsymbol{\theta}^*) - \mathbf{A}(\bar{\boldsymbol{\theta}})\| + \|\mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}^*) - \mathbf{A}(\bar{\boldsymbol{\theta}})\mathbf{B}(\boldsymbol{\theta}_0)\|\|\mathbf{A}(\bar{\boldsymbol{\theta}})\| \tag{93}$$

$$\leq A^*B^* \frac{\epsilon}{3 * \max(A^*B^*, A_{bnd}B_{bnd})} \tag{94}$$

$$+ \|\mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}^*) - \mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}_0) + \mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}_0) - \mathbf{A}(\bar{\boldsymbol{\theta}})\mathbf{B}(\boldsymbol{\theta}_0)\|\|\mathbf{A}(\bar{\boldsymbol{\theta}})\| \tag{95}$$

$$\leq \frac{\epsilon}{3} + \|\mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}^*) - \mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}_0)\|\|\mathbf{A}(\bar{\boldsymbol{\theta}})\| + \|\mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}_0) - \mathbf{A}(\bar{\boldsymbol{\theta}})\mathbf{B}(\boldsymbol{\theta}_0)\|\|\mathbf{A}(\bar{\boldsymbol{\theta}})\| \tag{96}$$

$$\leq \frac{\epsilon}{3} + \|\mathbf{A}(\boldsymbol{\theta}^*)\|\|\mathbf{B}(\boldsymbol{\theta}^*) - \mathbf{B}(\boldsymbol{\theta}_0)\|\|\mathbf{A}(\bar{\boldsymbol{\theta}})\| + \|\mathbf{A}(\boldsymbol{\theta}^*) - \mathbf{A}(\bar{\boldsymbol{\theta}})\|\|\mathbf{B}(-\theta_0)\|\|\mathbf{A}(\bar{\boldsymbol{\theta}})\| \tag{97}$$

$$\leq \frac{\epsilon}{3} + A^* \frac{\epsilon}{3A^*A_{bnd}} A_{bnd} + \frac{\epsilon}{3 * \max(A^*B^*, A_{bnd}B_{bnd})} B_{bnd}A_{bnd} \tag{98}$$

$$\leq \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} \tag{99}$$

$$= \epsilon \tag{100}$$

and the proof is done. □

Lemma 8. *With probability one, for all $\epsilon > 0$ there exists a $\delta > 0$ and an N such that for all $n > N$ and all $\boldsymbol{\theta}_0, \boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_k$ with $\|\boldsymbol{\theta}_i - \boldsymbol{\theta}^*\| < \delta$ for $i = 0, 1, \ldots, k$, $\|\mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\boldsymbol{\theta}_0)\mathbf{A}_n(\bar{\boldsymbol{\theta}}) - \mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}^*)\mathbf{A}(\boldsymbol{\theta}^*)\| < \epsilon$ for $\mathbf{A}_n(\bar{\boldsymbol{\theta}})$ defined above.*

Proof. We have by [Jennrich, 1969], Theorem 2, that the almost sure convergence of $\mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\boldsymbol{\theta}_0)\mathbf{A}_n(\bar{\boldsymbol{\theta}})$ to $\mathbf{A}(\bar{\boldsymbol{\theta}})\mathbf{B}(\boldsymbol{\theta}_0)\mathbf{A}(\bar{\boldsymbol{\theta}})$ is uniform in $\boldsymbol{\theta}$.

Consequently, pick $\epsilon > 0$. With probability one, there is an N such that for all $n > N$, and for all $\boldsymbol{\theta} \in \boldsymbol{\Theta}$, $\|\mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\boldsymbol{\theta}_0)\mathbf{A}_n(\bar{\boldsymbol{\theta}}) - \mathbf{A}(\bar{\boldsymbol{\theta}})\mathbf{B}(\boldsymbol{\theta}_0)\mathbf{A}(\bar{\boldsymbol{\theta}})\| < \epsilon/2$ with $\mathbf{A}_n(\bar{\boldsymbol{\theta}})$ as defined above.

For that same $\epsilon > 0$ by Lemma 7 there is a $\delta > 0$ such that for all $\boldsymbol{\theta}_0, \boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_k$ with $\|\boldsymbol{\theta}_i - \boldsymbol{\theta}^*\| < \delta$ for $i = 0, 1, \ldots, k$, $\|\mathbf{A}(\bar{\boldsymbol{\theta}})\mathbf{B}(\boldsymbol{\theta}_0)\mathbf{A}(\bar{\boldsymbol{\theta}}) - \mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}^*)\mathbf{A}(\boldsymbol{\theta}^*)\| < \epsilon/2$

Combining the two inequalities, we see that with probability one, for any $\epsilon > 0$ there is both $\delta > 0$ and N such that for all $n > N$ and for all $\boldsymbol{\theta}_0, \boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_k$ with $\|\boldsymbol{\theta}_i - \boldsymbol{\theta}^*\| < \delta$ for $i = 0, 1, \ldots, k$, we have $\|\mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\boldsymbol{\theta}_0)\mathbf{A}_n(\bar{\boldsymbol{\theta}}) - \mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}^*)\mathbf{A}(\boldsymbol{\theta}^*)\| < \epsilon$. □

Lemma 9. *Suppose* $\mathbf{A}$ *is a* $p \times p$ *matrix, and* $\mathbf{x}$ *is a* p*-dimensional vector. Then* $(-1)\|\mathbf{A}\| \cdot p^2 \cdot \|\mathbf{x}\|^2 \leq \mathbf{x}^T\mathbf{A}\mathbf{x} \leq \|\mathbf{A}\| \cdot p^2 \cdot \|\mathbf{x}\|^2$

Proof. We will prove the right-hand inequality. The left-hand inequality is done similarly.

$$\begin{aligned}
\mathbf{x}^T\mathbf{A}\mathbf{x} &= \sum_{i,j} a_{i,j}x_i x_j && (101)\\
&\leq |\sum_{i,j} a_{i,j}x_i x_j| && (102)\\
&\leq \sum_{i,j} |a_{i,j}x_i x_j| && (103)\\
&\leq \sum_{i,j} \max_{i,j}|a_{i,j}| \cdot |x_i x_j| && (104)\\
&= \max_{i,j}|a_{i,j}| \left(\sum_i |x_i|\right)^2 && (105)\\
&\leq \max_{i,j}|a_{i,j}| \cdot p^2 \cdot \max_i |x_i|^2 && (106)\\
&\leq \|\mathbf{A}\| \cdot p^2 \cdot \|\mathbf{x}\|^2 && (107)
\end{aligned}$$

□

We require that the eigenvalues of all relevant inverse sandwich matrix variations be finite and positive, and specifically bounded away from zero. First we look at $\mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}^*)^{-1}\mathbf{A}(\boldsymbol{\theta}^*)$. The standard regularity assumptions in [White, 1982] state that $\mathbf{A}(\boldsymbol{\theta}^*)$ is full rank and that $\mathbf{B}(\boldsymbol{\theta}^*)$ is invertible. That means that $\mathbf{B}(\boldsymbol{\theta}^*)$ is positive definite. Which implies that the asymptotic inverse sandwich, $\mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}^*)^{-1}\mathbf{A}(\boldsymbol{\theta}^*)$, is also positive definite. This means that the asymptotic inverse sandwich has eigenvalues that are all bounded and positive. In particular, there is an eigenvalue

with smallest magnitude, which we will call λ^*_{min} and an eigenvalue with largest magnitude, which we will call λ^*_{max}. For completeness, we have $0 < \lambda^*_{min} \leq \lambda^*_{max} < \infty$

Bounding the eigenvalues of $\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\mathbf{B}_n(\hat{\boldsymbol{\theta}}_n)^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)$ is now relatively easy. Since we have $\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\mathbf{B}_n(\hat{\boldsymbol{\theta}}_n)^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n) \rightarrow_{a.s.} \mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}^*)^{-1}\mathbf{A}(\boldsymbol{\theta}^*)$, we can bound the eigenvalues of $\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\mathbf{B}_n(\hat{\boldsymbol{\theta}}_n)^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)$ arbitrarily closely to the eigenvalues of $\mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}^*)^{-1}\mathbf{A}(\boldsymbol{\theta}^*)$. For any $\epsilon > 0$ with $\epsilon < \lambda^*_{min}$ we can then bound the range of eigenvalues of $\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\mathbf{B}_n(\hat{\boldsymbol{\theta}}_n)^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)$ to be between $\lambda^*_{min}-\epsilon$ and $\lambda^*_{max}+\epsilon$, for n sufficiently large enough.

To bound the eigenvalues of $\mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\boldsymbol{\theta})^{-1}\mathbf{A}_n(\bar{\boldsymbol{\theta}})$, we need another few steps. These are done with Lemmas 6, 7, and 8, which give us the results we need. Namely that with probability one, for all $\epsilon > 0$ there is both a $\delta > 0$ and N such that for all possible combinations of $\boldsymbol{\theta}$ in the ball of radius δ around $\boldsymbol{\theta}^*$ and for all $n > N$, $\|\mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\boldsymbol{\theta}_0)^{-1}\mathbf{A}_n(\bar{\boldsymbol{\theta}}) - \mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}^*)^{-1}\mathbf{A}(\boldsymbol{\theta}^*)\| < \epsilon$. This means that with a judicious choice of ϵ ($\epsilon = \lambda^*_{min}/2$ for example), we can bound the eigenvalues of $\mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\boldsymbol{\theta}_0)^{-1}\mathbf{A}_n(\bar{\boldsymbol{\theta}})$ to be positive and away from zero using Theorem 5.3 in [Zhan, 2013].

Lemma 10. *For any solution, $\breve{\boldsymbol{\theta}}_{n,PV} = \hat{\boldsymbol{\theta}}_n - \tilde{\boldsymbol{\theta}}_{n,PV}$, to the pivot equality, with probability one $\|\tilde{\boldsymbol{\theta}}_{n,PV}\| = O(1/\sqrt{n})$.*

Proof. Select ϵ such that $0 < \epsilon < \lambda^*_{min}/2$. With probability one, there is a $\delta > 0$ and an N_ϵ such that for all combinations of $\boldsymbol{\theta}_0, \boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_k$ with $\|\boldsymbol{\theta}_i - \boldsymbol{\theta}^*\| < \delta,\ i = 0, 1, \ldots, k$ and all $n > N$ we have $\|\mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\boldsymbol{\theta}_0)^{-1}\mathbf{A}_n(\bar{\boldsymbol{\theta}}) - \mathbf{A}(\boldsymbol{\theta}^*)\mathbf{B}(\boldsymbol{\theta}^*)^{-1}\mathbf{A}(\boldsymbol{\theta}^*)\| < \epsilon < \lambda^*_{min}/2$. This implies that the eigenvalues $\lambda_1, \ldots, \lambda_p$ of $\mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\boldsymbol{\theta}_0)^{-1}\mathbf{A}_n(\bar{\boldsymbol{\theta}})$ satisfy $0 < \lambda^*_{min}/2 < \lambda_j < \lambda^*_{max} + \lambda^*_{min}/2 < \infty,\ j = 1, \ldots, p$. Let us simplify this and say that there are bounds $0 < \lambda_{lo}$ and $\lambda_{hi} < \infty$ such that $0 < \lambda_{lo} < \lambda_j < \lambda_{hi} < \infty,\ j = 1, \ldots, p$.

By the almost sure convergence of $\hat{\boldsymbol{\theta}}_n$, we have with probability one that there exists an N_δ such that for all $n > N_\delta$, $\|\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*\| < \delta/2$, where δ is defined in the previous paragraph. This means that with probability one, the $\delta/2$ ball around $\hat{\boldsymbol{\theta}}_n$ is contained within the δ ball around $\boldsymbol{\theta}^*$. We can take $N_{max} = \max(N_\epsilon, N_\delta)$ so that all of the results in this and the previous paragraph hold.

Consider the case $n > N_{max}$, and select a point on the boundary of the $\delta/2$ ball around $\hat{\boldsymbol{\theta}}_n$. Let's refer to this point as $\boldsymbol{\theta}_0 = \hat{\boldsymbol{\theta}}_n - \tilde{\boldsymbol{\theta}}_n$ because we want to emphasize the fact that $\|\tilde{\boldsymbol{\theta}}_n\| = \delta/2$. We can plug this point in for $\boldsymbol{\theta}$ in the pivotal quadratic form. Recall we also use it and $\hat{\boldsymbol{\theta}}_n$ and as

the points to define $\mathbf{A}_n(\bar{\boldsymbol{\theta}})$. Using results on the bounds of quadratic forms of symmetric positive definite matrices we have

$$\lambda_{lo}\|\tilde{\boldsymbol{\theta}}_n\|^2 \leq (\tilde{\boldsymbol{\theta}}_n)^T\mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\boldsymbol{\theta}_0)^{-1}\mathbf{A}_n(\bar{\boldsymbol{\theta}})(\tilde{\boldsymbol{\theta}}_n) \leq \lambda_{hi}\|\tilde{\boldsymbol{\theta}}_n\|^2 \tag{108}$$

$$\lambda_{lo}\frac{\delta^2}{4} \leq (\tilde{\boldsymbol{\theta}}_n)^T\mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\boldsymbol{\theta}_0)^{-1}\mathbf{A}_n(\bar{\boldsymbol{\theta}})(\tilde{\boldsymbol{\theta}}_n) \leq \lambda_{hi}\frac{\delta^2}{4}. \tag{109}$$

which is true regardless of the direction of $\tilde{\boldsymbol{\theta}}_n$, hence it is uniformly true in all directions away from $\hat{\boldsymbol{\theta}}_n$.

We now compare (109) with the main pivot inequality in (53). As n increases, we see that the lower bound in (109) is a constant. However as n increases in (53), it will eventually be the case that $\chi^2_{p,1-\alpha}/n < \lambda_{lo}\frac{\delta^2}{4}$, namely when $n > \frac{4\chi^2_{p,1-\alpha}}{\lambda_{lo}\delta^2}$. Since $\hat{\boldsymbol{\theta}}_n$ makes the pivotal inequality equal to zero, and every point on the $\delta/2$ ball makes the quadratic form evaluate to more than $\chi^2_{p,1-\alpha}/n$, we can invoke the intermediate value theorem to conclude that all solutions to the pivotal inequality will lie within the $\delta/2$ ball around $\hat{\boldsymbol{\theta}}_n$, and that they do exist. We can now take $N = \max(\frac{4\chi^2_{p,1-\alpha}}{\lambda_{lo}\delta^2}, N_{max})$.

Define N and δ as we have done. We consider a solution $\breve{\boldsymbol{\theta}}_{n,PV} = \hat{\boldsymbol{\theta}}_n - \tilde{\boldsymbol{\theta}}_{n,PV}$ to the pivotal equality in an arbitrary direction from $\hat{\boldsymbol{\theta}}$, and specifically look at the equality

$$(\tilde{\boldsymbol{\theta}}_{n,PV})^T\mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\breve{\boldsymbol{\theta}}_{n,PV})^{-1}\mathbf{A}_n(\bar{\boldsymbol{\theta}})(\tilde{\boldsymbol{\theta}}_{n,PV}) = \frac{\chi^2_{p,1-\alpha}}{n}. \tag{110}$$

Given the eigenvalue-based bounds above and using knowledge of quadratic forms, we can conclude

$$\lambda_{lo}\|\tilde{\boldsymbol{\theta}}_{n,PV}\|^2 \leq (\tilde{\boldsymbol{\theta}}_{n,PV})^T\mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\breve{\boldsymbol{\theta}}_{n,PV})^{-1}\mathbf{A}_n(\bar{\boldsymbol{\theta}})(\tilde{\boldsymbol{\theta}}_{n,PV}) \leq \lambda_{hi}\|\tilde{\boldsymbol{\theta}}_{n,PV}\|^2 \tag{111}$$

$$\lambda_{lo}\|\tilde{\boldsymbol{\theta}}_{n,PV}\|^2 \leq \frac{\chi^2_{p,1-\alpha}}{n} \leq \lambda_{hi}\|\tilde{\boldsymbol{\theta}}_{n,PV}\|^2 \tag{112}$$

and solving for $\|\tilde{\boldsymbol{\theta}}_{n,PV}\|$ we have

$$\sqrt{\frac{\chi^2_{p,1-\alpha}}{n\lambda_{hi}}} \leq \|\tilde{\boldsymbol{\theta}}_{n,PV}\| \leq \sqrt{\frac{\chi^2_{p,1-\alpha}}{n\lambda_{lo}}} \tag{113}$$

which gives us the result that $\|\tilde{\boldsymbol{\theta}}_{n,PV}\| = O(1/\sqrt{n})$ with probability one. □

Proof of Theorem 1

Proof. Consider solutions $\breve{\boldsymbol{\theta}}_{n,PV}$ and $\breve{\boldsymbol{\theta}}_{n,SW}$ to the pivot equality and sandwich equality, respectively and both are in the same direction from $\hat{\boldsymbol{\theta}}_n$. First, rewrite them as

$$\breve{\boldsymbol{\theta}}_{n,PV} = \hat{\boldsymbol{\theta}}_n - \tilde{\boldsymbol{\theta}}_{n,PV} \tag{114}$$

$$\breve{\boldsymbol{\theta}}_{n,SW} = \hat{\boldsymbol{\theta}}_n - \tilde{\boldsymbol{\theta}}_{n,SW} \tag{115}$$

where we emphasize again that $\tilde{\boldsymbol{\theta}}_{n,PV}$ and $\tilde{\boldsymbol{\theta}}_{n,SW}$ are vectors having the same direction, but with possibly different magnitudes.

We have that $\|\tilde{\boldsymbol{\theta}}_{n,PV}\|$ is $O(1/\sqrt{n})$ with probability one by Lemma 10. This also tells us that with probability one as $n \to \infty$, $\breve{\boldsymbol{\theta}}_{n,PV}$ will become arbitrarily close to $\hat{\boldsymbol{\theta}}_n$ and specifically, there is a constant A such that for sufficiently large n, $\|\tilde{\boldsymbol{\theta}}_{n,PV}\| \leq A/\sqrt{n}$. We also have that $\|\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\mathbf{B}_n(\hat{\boldsymbol{\theta}}_n)^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n) - \mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\breve{\boldsymbol{\theta}}_{n,PV})^{-1}\mathbf{A}_n(\bar{\boldsymbol{\theta}})\| \to_{a.s.} 0$ by the continuity of the matrix functions $\mathbf{A}_n(\boldsymbol{\theta})$ and $\mathbf{B}_n(\boldsymbol{\theta})^{-1}$. Pick $\zeta > 0$ small enough so that $\zeta < \chi^2_{p,1-\alpha}$. With probability one there is an M such that for all $n > M$, $\|\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\mathbf{B}_n(\hat{\boldsymbol{\theta}}_n)^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n) - \mathbf{A}_n(\bar{\boldsymbol{\theta}}_n)\mathbf{B}_n(\breve{\boldsymbol{\theta}}_{n,PV})^{-1}\mathbf{A}_n(\bar{\boldsymbol{\theta}}_n)\| < \frac{\zeta}{p^2A^2}$.

Now take a solution, $\breve{\boldsymbol{\theta}}_{n,PV} = \hat{\boldsymbol{\theta}}_n - \tilde{\boldsymbol{\theta}}_{n,PV}$, to the pivot equality and plug it into the sandwich equality

$$(\hat{\boldsymbol{\theta}}_n - \breve{\boldsymbol{\theta}}_{n,PV})^T\mathbf{A}_n(\hat{\boldsymbol{\theta}})\mathbf{B}_n(\hat{\boldsymbol{\theta}})^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}})(\hat{\boldsymbol{\theta}}_n - \breve{\boldsymbol{\theta}}_{n,PV}) = \tag{116}$$

$$(\tilde{\boldsymbol{\theta}}_{n,PV})^T\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\mathbf{B}_n(\hat{\boldsymbol{\theta}}_n)^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)(\tilde{\boldsymbol{\theta}}_{n,PV}) = \tag{117}$$

$$(\tilde{\boldsymbol{\theta}}_{n,PV})^T[\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\mathbf{B}_n(\hat{\boldsymbol{\theta}}_n)^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n) - \tag{118}$$

$$\mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\breve{\boldsymbol{\theta}}_{n,PV})^{-1}\mathbf{A}_n(\bar{\boldsymbol{\theta}}) + \mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\breve{\boldsymbol{\theta}}_{n,PV})^{-1}\mathbf{A}_n(\bar{\boldsymbol{\theta}})](\tilde{\boldsymbol{\theta}}_{n,PV}) = \tag{119}$$

$$(\tilde{\boldsymbol{\theta}}_{n,PV})^T\mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\breve{\boldsymbol{\theta}}_{n,PV})^{-1}\mathbf{A}_n(\bar{\boldsymbol{\theta}})(\tilde{\boldsymbol{\theta}}_{n,PV}) + \tag{120}$$

$$(\tilde{\boldsymbol{\theta}}_{n,PV})^T[\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\mathbf{B}_n(\hat{\boldsymbol{\theta}}_n)^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n) - \mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\breve{\boldsymbol{\theta}}_{n,PV})^{-1}\mathbf{A}_n(\bar{\boldsymbol{\theta}})](\tilde{\boldsymbol{\theta}}_{n,PV}) = \tag{121}$$

$$\frac{\chi^2_{p,1-\alpha}}{n} + (\tilde{\boldsymbol{\theta}}_{n,PV})^T[\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\mathbf{B}_n(\hat{\boldsymbol{\theta}}_n)^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n) - \mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\breve{\boldsymbol{\theta}}_{n,PV})^{-1}\mathbf{A}_n(\bar{\boldsymbol{\theta}})](\tilde{\boldsymbol{\theta}}_{n,PV}). \tag{122}$$

By the previous part of this paragraph we have that

$$
\begin{aligned}
-\frac{A^2}{n}p^2\frac{\zeta}{p^2A^2} \leq (\tilde{\boldsymbol{\theta}}_{n,PV})^T \Big[&\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\mathbf{B}_n(\hat{\boldsymbol{\theta}}_n)^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n) \\
&-\mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\breve{\boldsymbol{\theta}}_{n,PV})^{-1}\mathbf{A}_n(\bar{\boldsymbol{\theta}})\Big](\tilde{\boldsymbol{\theta}}_{n,PV}) \leq \frac{A^2}{n}p^2\frac{\zeta}{p^2A^2} \\
-\frac{\zeta}{n} \leq (\tilde{\boldsymbol{\theta}}_{n,PV})^T[&\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\mathbf{B}_n(\hat{\boldsymbol{\theta}}_n)^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n) - \mathbf{A}_n(\bar{\boldsymbol{\theta}})\mathbf{B}_n(\breve{\boldsymbol{\theta}}_{n,PV})^{-1}\mathbf{A}_n(\bar{\boldsymbol{\theta}})](\tilde{\boldsymbol{\theta}}_{n,PV}) \leq \frac{\zeta}{n}
\end{aligned} \tag{123}
$$

and plugging Line (123) into Line (122) we have

$$
\frac{\chi^2_{p,1-\alpha}-\zeta}{n} \leq (\tilde{\boldsymbol{\theta}}_{n,PV})^T\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\mathbf{B}_n(\hat{\boldsymbol{\theta}})^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)(\tilde{\boldsymbol{\theta}}_{n,PV}) \leq \frac{\chi^2_{p,1-\alpha}+\zeta}{n} \tag{124}
$$

Taking each side of the above inequality separately, we see

$$
\frac{\chi^2_{p,1-\alpha}}{\chi^2_{p,1-\alpha}+\zeta}(\tilde{\boldsymbol{\theta}}_{n,PV})^T\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\mathbf{B}_n(\hat{\boldsymbol{\theta}}_n)^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)(\tilde{\boldsymbol{\theta}}_{n,PV}) < \frac{\chi^2_{p,1-\alpha}}{n} \tag{125}
$$

$$
\frac{\chi^2_{p,1-\alpha}}{\chi^2_{p,1-\alpha}-\zeta}(\tilde{\boldsymbol{\theta}}_{n,PV})^T\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\mathbf{B}_n(\hat{\boldsymbol{\theta}}_n)^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)(\tilde{\boldsymbol{\theta}}_{n,PV}) > \frac{\chi^2_{p,1-\alpha}}{n} \tag{126}
$$

Now we can rewrite $\frac{\chi^2_{p,1-\alpha}}{\chi^2_{p,1-\alpha}+\zeta}$ as $1-\frac{\zeta}{\chi^2_{p,1-\alpha}+\zeta} = 1-\xi_+$. Similarly, $\frac{\chi^2_{p,1-\alpha}}{\chi^2_{p,1-\alpha}-\zeta} = 1+\frac{\zeta}{\chi^2_{p,1-\alpha}-\zeta} = 1+\xi_-$. Note in particular that ζ and ξ_- are in one-to-one correspondence, and as $\zeta \to 0$, $\xi_- \to 0$. This gives us

$$
(1-\xi_+)(\tilde{\boldsymbol{\theta}}_{n,PV})^T\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\mathbf{B}_n(\hat{\boldsymbol{\theta}}_n)^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)(\tilde{\boldsymbol{\theta}}_{n,PV}) < \frac{\chi^2_{p,1-\alpha}}{n} \tag{127}
$$

$$
(1+\xi_-)(\tilde{\boldsymbol{\theta}}_{n,PV})^T\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\mathbf{B}_n(\hat{\boldsymbol{\theta}}_n)^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)(\tilde{\boldsymbol{\theta}}_{n,PV}) > \frac{\chi^2_{p,1-\alpha}}{n}. \tag{128}
$$

Since $\xi_- > \xi_+$ and $1-\xi_- < 1-\xi_+$ we can rewrite the first inequality to get

$$
(1-\xi_-)(\tilde{\boldsymbol{\theta}}_{n,PV})^T\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\mathbf{B}_n(\hat{\boldsymbol{\theta}}_n)^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)(\tilde{\boldsymbol{\theta}}_{n,PV}) < \frac{\chi^2_{p,1-\alpha}}{n} \tag{129}
$$

$$
(1+\xi_-)(\tilde{\boldsymbol{\theta}}_{n,PV})^T\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)\mathbf{B}_n(\hat{\boldsymbol{\theta}}_n)^{-1}\mathbf{A}_n(\hat{\boldsymbol{\theta}}_n)(\tilde{\boldsymbol{\theta}}_{n,PV}) > \frac{\chi^2_{p,1-\alpha}}{n}. \tag{130}
$$

We have that $\hat{\boldsymbol{\theta}}_n - (\sqrt{1-\xi_-})\tilde{\boldsymbol{\theta}}_{n,PV}$ satisfies the plug-in sandwich quadratic form inequality but $\hat{\boldsymbol{\theta}}_n - (\sqrt{1+\xi_-})\tilde{\boldsymbol{\theta}}_{n,PV}$ does not. By the continuity of the plug-in sandwich quadratic form, the solution to the sandwich equality, i.e. $\breve{\boldsymbol{\theta}}_{n,SW} = \hat{\boldsymbol{\theta}}_n - \tilde{\boldsymbol{\theta}}_{n,SW}$, must lie on the line segment between these two points.

Once we have selected ζ and consequently ξ_- as described above, $1 - \sqrt{1-\xi_-} > \sqrt{1+\xi_-} - 1$. This means that $\hat{\boldsymbol{\theta}}_n - (\sqrt{1-\xi_-})\tilde{\boldsymbol{\theta}}_{n,PV}$ is the furthest point from $\hat{\boldsymbol{\theta}}_n - \tilde{\boldsymbol{\theta}}_{n,PV}$ in the interval $(\hat{\boldsymbol{\theta}}_n - (\sqrt{1-\xi_-})\tilde{\boldsymbol{\theta}}_{n,PV}, \hat{\boldsymbol{\theta}}_n - (\sqrt{1+\xi_-})\tilde{\boldsymbol{\theta}}_{n,PV})$. Putting this all together, we have that for all $1 > \xi_- > 0$ with probability one there is an M such that for all $n > M$

$$\begin{aligned}
\sqrt{n}\|(\hat{\boldsymbol{\theta}}_n - \tilde{\boldsymbol{\theta}}_{n,PV}) - (\hat{\boldsymbol{\theta}}_n - \tilde{\boldsymbol{\theta}}_{n,SW})\| &= \sqrt{n}\|\tilde{\boldsymbol{\theta}}_{n,SW} - \tilde{\boldsymbol{\theta}}_{n,PV}\| && (131)\\
&\leq \|(\sqrt{1-\xi_-})\tilde{\boldsymbol{\theta}}_{n,PV} - \tilde{\boldsymbol{\theta}}_{n,PV}\| && (132)\\
&= \sqrt{n}\|\tilde{\boldsymbol{\theta}}_{n,PV}\| \mid 1 - \sqrt{1-\xi_-} \mid && (133)\\
&\leq \sqrt{n}\sqrt{\frac{\chi^2_{p,1-\alpha}}{n\lambda_{lo}}}(\mid 1 - \sqrt{1-\xi_-} \mid) && (134)\\
&= \sqrt{\frac{\chi^2_{p,1-\alpha}}{\lambda_{lo}}}(\mid 1 - \sqrt{1-\xi_-} \mid) && (135)
\end{aligned}$$

but ζ was chosen to be arbitrarily small, which gives us that ξ_- is also arbitrarily small. Hence we have that $\sqrt{n}\|(\hat{\boldsymbol{\theta}}_n - \tilde{\boldsymbol{\theta}}_{n,PV}) - (\hat{\boldsymbol{\theta}}_n - \tilde{\boldsymbol{\theta}}_{n,SW})\| \rightarrow_{a.s.} 0$ as $n \rightarrow \infty$, as desired. □

References

H. Akaike. Information theory and an extension of the maximum likelihood principle. *IEEE International Symposium on Information Theory, Proceedings*, pages 267–281, 1973.

Takeshi Amemiya. Regression analysis when the dependent variable is truncated normal. *Econometrica*, 41(6):997–1016, 1973.

Norman Breslow. Tests of hypotheses in overdispersed poisson regression and other quasi-likelihood models. *Journal of the American Statistical Association*, 85(410):565–571, 1990.

Xiaohong Chen and Yanqin Fan. Pseudo-likelihood ratio tests for semiparametric multivariate copula model selection. *Canadian Journal of Statistics*, 33(3):389–414, 2005.

Peter Diggle, Patrick Heagerty, Kung-Yee Liang, and Scott Zeger. *Analysis of longitudinal data, 2nd ed.* Oxford University Press, 2009.

B Efron. Discussion: jackknife, bootstrap and other resampling methods in regression analysis. *The Annals of Statistics*, pages 1301–1304, 1986.

Bradley Efron. *The jackknife, the bootstrap and other resampling plans*, volume 38. SIAM, 1982.

Michael P Fay and Barry I Graubard. Small-sample adjustments for wald-type tests using sandwich estimators. *Biometrics*, 57(4):1198–1206, 2001.

D. Firth. Discussion of the paper by liang, zeger & qaqish multivariate regression analysis for categorical data. *Journal of the Royal Statistical Society, Series B*, 54:24–26, 1992.

M Furi and M Martelli. On the mean value theorem, inequality, and inclusion. *The American Mathematical Monthly*, 98(9):840–846, 1991.

James W Hardin. The sandwich estimate of variance. *Advances in Econometrics*, 17:45–73, 2003.

David V Hinkley. Jackknifing in unbalanced situations. *Technometrics*, 19(3):285–292, 1977.

Susan D Horn, Roger A Horn, and David B Duncan. Estimating heteroscedastic variances in linear models. *Journal of the American Statistical Association*, 70(350):380–385, 1975.

Peter J Huber. The behavior of maximum likelihood estimates under nonstandard conditions. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 221–233, 1967.

Peter J Huber and Elvezio M Ronchetti. *Robust statistics*. Wiley, 2009.

Robert I Jennrich. Asymptotic properties of non-linear least squares estimators. *The Annals of Mathematical Statistics*, 40(2):633–643, 1969.

Göran Kauermann and Raymond J Carroll. The sandwich variance estimator: efficiency properties and coverage probability of confidence intervals. 2000.

Göran Kauermann and Raymond J Carroll. A note on the efficiency of sandwich covariance matrix estimation. *Journal of the American Statistical Association*, 96(456):1387–1396, 2001.

Erich Leo Lehmann. *Elements of large-sample theory*. Springer Science & Business Media, 1999.

Kung-Yee Liang, Scott L Zeger, and Bahjat Qaqish. Multivariate regression analyses for categorical data. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 3–40, 1992.

Peter McCullagh. Discussion of the paper by liang, zeger & qaqish multivariate regression analysis for categorical data. *Journal of the Royal Statistical Society, Series B*, 54:33–34, 1992.

John F Monahan and Dennis D Boos. Proper likelihoods for bayesian analysis. *Biometrika*, 79(2): 271–278, 1992.

Halbert White. Maximum likelihood estimation of misspecified models. *Econometrica: Journal of the Econometric Society*, pages 1–25, 1982.

Chien-Fu Jeff Wu. Jackknife, bootstrap and other resampling methods in regression analysis. *the Annals of Statistics*, pages 1261–1295, 1986.

Xingzhi Zhan. *Matrix theory*, volume 147. American Mathematical Soc., 2013.